

Sparse sum-of-squares certificates on finite abelian groups

Hamza Fawzi

James Saunderson

Pablo A. Parrilo

March 3, 2015

Abstract

Let G be a finite abelian group. This paper is concerned with nonnegative functions on G that are *sparse* with respect to the Fourier basis. We establish combinatorial conditions on subsets \mathcal{S} and \mathcal{T} of Fourier basis elements under which nonnegative functions with Fourier support \mathcal{S} are sums of squares of functions with Fourier support \mathcal{T} . Our combinatorial condition involves constructing a chordal cover of a graph related to G and \mathcal{S} (the Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$) with maximal cliques related to \mathcal{T} . Our result relies on two main ingredients: the decomposition of sparse positive semidefinite matrices with a chordal sparsity pattern, as well as a simple but key observation exploiting the structure of the Fourier basis elements of G (the characters of G).

We apply our general result to two examples. First, in the case where $G = \mathbb{Z}_2^n$, by constructing a particular chordal cover of the half-cube graph, we prove that any nonnegative quadratic form in n binary variables is a sum of squares of functions of degree at most $\lceil n/2 \rceil$, establishing a conjecture of Laurent. Second, we consider nonnegative functions of degree d on \mathbb{Z}_N (when d divides N). By constructing a particular chordal cover of the d th power of the N -cycle, we prove that any such function is a sum of squares of functions with at most $3d \log(N/d)$ nonzero Fourier coefficients. Dually this shows that a certain cyclic polytope in \mathbb{R}^{2d} with N vertices can be expressed as a projection of a section of the cone of positive semidefinite matrices of size $3d \log(N/d)$. Putting $N = d^2$ gives a family of polytopes in \mathbb{R}^{2d} with linear programming extension complexity $\Omega(d^2)$ and semidefinite programming extension complexity $O(d \log(d))$. To the best of our knowledge, this is the first explicit family of polytopes (P_d) in increasing dimensions where $\text{xcpsd}(P_d) = o(\text{xclp}(P_d))$ (where xcpsd and xclp are respectively the SDP and LP extension complexity).

The authors are with the Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. Email: {hfawzi,jamess,parrilo}@mit.edu.

Contents

1	Introduction	2
2	Preliminaries	6
2.1	Fourier analysis on finite groups	6
2.2	Chordal graphs and matrix completion	8
2.3	Lifts of polytopes	9
3	Main result for general finite abelian groups	9
3.1	Nonnegative functions and sum-of-squares certificates	10
3.2	Dual point of view: moment matrices and matrix completion	13
3.2.1	Moment polytopes	13
3.2.2	Moment polytopes and matrix completion	15
3.3	Real sums-of-squares and moment polytopes	16
4	Application 1: The case $G = \mathbb{Z}_2^n$ and binary quadratic optimization	18
4.1	Quadratic forms on $\{-1, 1\}^n$ and the cut polytope	18
4.2	The associated Cayley graph	19
4.3	Applying Theorem 1	20
5	Application 2: The case $G = \mathbb{Z}_N$ and cyclic polytopes	21
5.1	The case $\mathcal{S} = \{-1, 0, 1\}$: the cycle graph	21
5.2	Degree d functions: powers of cycle graph	21
5.2.1	Triangulating the Cayley graph	22
5.2.2	Cyclic polytopes	24
6	Conclusion	26
A	Additional proofs	26
A.1	Proofs from Section 3.2	26
A.2	Proof of Lemma 1	28
A.3	Proof of Proposition 6	28
A.4	Proof of Theorem 9: triangulation of the cycle graph	29

1 Introduction

Let G be a finite abelian group. It is well-known that any function $f : G \rightarrow \mathbb{C}$ admits a Fourier decomposition where the Fourier basis consists of the *characters* of G . Such a decomposition takes the form

$$f(x) = \sum_{\chi \in \widehat{G}} \widehat{f}(\chi) \chi(x) \quad \forall x \in G$$

where \widehat{G} is the set of characters of G (known as the *dual group* of G) and $\widehat{f}(\chi)$ are the Fourier coefficients of f . The function $f : G \rightarrow \mathbb{C}$ is called *sparse* if only a few of its Fourier coefficients are nonzero. More precisely we say that f is supported on $\mathcal{S} \subseteq \widehat{G}$ if $\widehat{f}(\chi) = 0$ whenever $\chi \notin \mathcal{S}$.

This paper is concerned with functions $f : G \rightarrow \mathbb{C}$ that are *sparse* and *nonnegative*, i.e., $f(x) \in \mathbb{R}_+$ for all $x \in G$. If f is a nonnegative function on G , a *sum-of-squares certificate* for the nonnegativity of f has the form:

$$f(x) = \sum_{j=1}^J |f_j(x)|^2 \quad \forall x \in G \tag{1}$$

where $f_j : G \rightarrow \mathbb{C}$. Sum-of-squares certificates of nonnegative functions play an important role in optimization and particularly in semidefinite programming [BPT13]. When the function f is sparse, it is natural

to ask whether f admits a sum-of-squares certificate that is also sparse, i.e., where all the functions f_j are supported on a common “small” set $\mathcal{T} \subseteq \widehat{G}$. This is the main question of interest in this paper:

$$\text{Given } \mathcal{S} \subseteq \widehat{G}, \text{ find a subset } \mathcal{T} \subseteq \widehat{G} \text{ such that any nonnegative function } G \rightarrow \mathbb{R}_+ \text{ supported on } \mathcal{S} \text{ admits a sum-of-squares certificate supported on } \mathcal{T}. \quad (\text{Q})$$

Our main result is to give a sufficient condition for a set \mathcal{T} to satisfy the requirement above for a given \mathcal{S} . The condition is expressed in terms of *chordal covers* of the Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$. Recall that the Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$ is the graph where nodes correspond to elements of \widehat{G} and where χ, χ' are connected by an edge if $\chi^{-1}\chi' \in \mathcal{S}$. Our main result can be stated as follows:

Theorem 1. *Let $\mathcal{S} \subseteq \widehat{G}$. Let \mathcal{T} be a subset of \widehat{G} obtained as follows: Let Γ be a chordal cover of $\text{Cay}(\widehat{G}, \mathcal{S})$, and for each maximal clique \mathcal{C} of Γ , let $\chi_{\mathcal{C}}$ be an element of \widehat{G} ; define*

$$\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\}) = \bigcup_{\mathcal{C}} \chi_{\mathcal{C}} \mathcal{C} \quad (2)$$

where the union is over all the maximal cliques of Γ and where $\chi_{\mathcal{C}} \mathcal{C} := \{\chi_{\mathcal{C}} \chi : \chi \in \mathcal{C}\}$ is the translation of \mathcal{C} by $\chi_{\mathcal{C}}$. Then any nonnegative function supported on \mathcal{S} admits a sum-of-squares certificate supported on $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$.

Theorem 1 gives a way to construct a set \mathcal{T} that satisfies the condition in (Q) for a given $\mathcal{S} \subseteq \widehat{G}$. Such a construction proceeds in two steps: first choose a chordal cover Γ of the graph $\text{Cay}(\widehat{G}, \mathcal{S})$, and then choose elements $\chi_{\mathcal{C}} \in \widehat{G}$ for each maximal clique \mathcal{C} of \widehat{G} . Different choices of Γ and $\{\chi_{\mathcal{C}}\}$ will in general lead to different sets $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$. When using Theorem 1, one wants to find a good choice of Γ and $\{\chi_{\mathcal{C}}\}$ such that the resulting set $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$ is as small as possible (or has other desirable properties).

One of the main strengths of Theorem 1 is in the ability to choose the elements $\{\chi_{\mathcal{C}}\}$. In fact the conclusion of Theorem 1 is almost trivial if $\chi_{\mathcal{C}} = 1_{\widehat{G}}$ for all \mathcal{C} , since in this case it simply says that any nonnegative function has a sum-of-squares certificate supported on \widehat{G} , which is easy to see since G is finite. As we will see in the applications, it is the ability to *translate* the cliques \mathcal{C} of Γ via the choice of $\chi_{\mathcal{C}}$ that is key in Theorem 1 and allows us to obtain interesting results. Equation (2) gives us the intuition behind a good choice of $\{\chi_{\mathcal{C}}\}$: in order to minimize the cardinality of $\mathcal{T}(\Gamma, \{\chi_{\mathcal{C}}\})$ one would like to find the translations $\chi_{\mathcal{C}}$ that maximize the total overlap of the cliques (i.e., minimize the cardinality of their union).

Before describing the main idea behind Theorem 1 and its proof, we illustrate how one can use Theorem 1 in two important special cases, namely $G = \mathbb{Z}_2^n$ (the boolean hypercube) and $G = \mathbb{Z}_N$.

- **Boolean hypercube:** Consider the case $G = \{-1, 1\}^n \cong \mathbb{Z}_2^n$. The Fourier expansion of functions on $\{-1, 1\}^n$ take the form

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i. \quad (3)$$

A function f is said to have degree d if $\widehat{f}(S) = 0$ for all S such that $|S| > d$. Many combinatorial optimization problems correspond to optimizing a certain function f over $\{-1, 1\}^n$. For example the maximum cut problem in graph theory consists in optimizing a *quadratic function* over $\{-1, 1\}^n$. In [Lau03] Laurent conjectured that any nonnegative quadratic function on the hypercube is a sum of squares of functions of degree at most $\lceil n/2 \rceil$. Using our notations, this corresponds to asking whether for $\mathcal{S} = \{S \subseteq [n] : |S| = 0 \text{ or } 2\}$ one can find $\mathcal{T} \subseteq \{S \subseteq [n] : |S| \leq \lceil n/2 \rceil\}$ such that the conclusion of Theorem 1 holds. By studying chordal covers of the Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$ we are able to answer this question positively:

Theorem 2. *Any nonnegative quadratic function on $\{-1, 1\}^n$ is a sum-of-squares of polynomials of degree at most $\lceil n/2 \rceil$.*

Note that Blekherman et al. [BGP14] previously showed a weaker version of the conjecture that allows for multipliers: They showed that for any nonnegative quadratic function f on the hypercube, there exists h sum-of-squares such that $h(x)f(x)$ is a sum-of-squares of polynomials of degree at most $\lceil n/2 \rceil$.

- **Trigonometric polynomials:** Another important application that we consider in this paper is the case where $G = \mathbb{Z}_N$, the (additive) group of integers modulo N . The Fourier decomposition of a function $f : \mathbb{Z}_N \rightarrow \mathbb{C}$ is the usual discrete Fourier transform and takes the form:

$$f(x) = \sum_{k \in \mathbb{Z}_N} \hat{f}(k) e^{2i\pi kx/N} \quad (4)$$

where $\hat{f}(k)$ are the Fourier coefficients of f . A function f is said to have degree d if $\text{supp } f \subseteq \{-d, -(d-1), \dots, d-1, d\}$. Nonnegative trigonometric polynomials play an important role in many areas such as in signal processing [Dum07], but also in convex geometry [Zie95, Bar02], in their relation to (trigonometric) cyclic polytopes. We are interested in nonnegative functions on G of degree at most d , i.e., functions supported on $\mathcal{S} = \{-d, -(d-1), \dots, d-1, d\}$. By studying chordal covers of $\text{Cay}(\widehat{G}, \mathcal{S})$ (which is nothing but the d 'th power of the cycle graph) and using Theorem 1 we are able to show the following:

Theorem 3. *Let N and d be two integers and assume that d divides N . Then there exists $\mathcal{T} \subseteq \mathbb{Z}_N$ with $|\mathcal{T}| \leq 3d \log(N/d)$ such that any nonnegative function on \mathbb{Z}_N of degree at most d has a sum-of-squares certificate supported on \mathcal{T} .*

Remark. Note that if one is interested in functions of degree at most d on \mathbb{Z}_N and d does not divide N , then one can still apply Theorem 3 with d' instead of d , where d' is the smallest divisor of N that is greater than d .

Dual point of view and moment polytopes Theorem 1 can be interpreted from the dual point of view as giving a semidefinite programming description of certain moment polytopes. If $\mathcal{S} \subseteq \widehat{G}$, define the *moment polytope* $\mathcal{M}(G, \mathcal{S})$ to be the set of \mathcal{S} -moments of probability distributions on G , i.e.,

$$\mathcal{M}(G, \mathcal{S}) = \left\{ (\mathbb{E}_{x \sim \mu} [\chi(x)])_{\chi \in \mathcal{S}} \in \mathbb{C}^{\mathcal{S}} : \mu \text{ a probability measure supported on } G \right\}.$$

Note that $\mathcal{M}(G, \mathcal{S})$ is a polytope since it can be equivalently expressed as:

$$\mathcal{M}(G, \mathcal{S}) = \text{conv} \left\{ (\chi(x))_{\chi \in \mathcal{S}} \in \mathbb{C}^{\mathcal{S}} : x \in G \right\}.$$

Note that from a geometric point of view, nonnegative functions $f : G \rightarrow \mathbb{R}_+$ supported on \mathcal{S} correspond to valid linear inequalities for the polytope $\mathcal{M}(G, \mathcal{S})$. By giving a sum-of-squares characterization for all valid inequalities of $\mathcal{M}(G, \mathcal{S})$ Theorem 1 allows us to obtain a semidefinite programming description of $\mathcal{M}(G, \mathcal{S})$. The following statement can be obtained from Theorem 1 by duality (we call this result “Theorem 1D” to reflect that it is a dual version of “Theorem 1”—we adopt this numbering convention throughout the paper):

Theorem 1D. *Let $\mathcal{S} \subseteq \widehat{G}$ and let $\mathcal{T} = \mathcal{T}(\Gamma, \{\chi_C\})$ be as defined in Theorem 1. Then we have the following semidefinite programming description of the moment polytope $\mathcal{M}(G, \mathcal{S})$:*

$$\begin{aligned} \mathcal{M}(G, \mathcal{S}) = \left\{ (\ell_\chi)_{\chi \in \mathcal{S}} : \exists (y_\chi)_{\chi \in \mathcal{T}^{-1}\mathcal{T}} \text{ such that } y_\chi = \ell_\chi \text{ for all } \chi \in \mathcal{S}, \text{ and} \right. \\ \left. y_{1_{\widehat{G}}} = 1, \text{ and } [y_{\overline{\chi\chi'}}]_{\chi, \chi' \in \mathcal{T}} \succeq 0 \right\}. \end{aligned} \quad (5)$$

In terms of positive semidefinite lifts, Equation (5) shows that $\mathcal{M}(G, \mathcal{S})$ has a *Hermitian positive semidefinite lift* of size $|\mathcal{T}|$. We now illustrate this dual point of view for the two applications mentioned above, $G = \{-1, 1\}^n$ and $G = \mathbb{Z}_N$:

- For the case of the boolean hypercube $G = \{-1, 1\}^n$, if $\mathcal{S} = \{S \subseteq [n] : |S| = 0 \text{ or } 2\}$, the moment polytope $\mathcal{M}(\{-1, 1\}^n, \mathcal{S} \setminus \{\emptyset\})$ is nothing but the *cut polytope* for the complete graph on n vertices which we denote by CUT_n :

$$\text{CUT}_n = \text{conv} \left\{ (x_i x_j)_{i < j} \in \mathbb{R}^{\binom{n}{2}} : x \in \{-1, 1\}^n \right\}.$$

From the dual point of view, Theorem 2 shows that the $\lceil n/2 \rceil$ level of the Lasserre hierarchy for the cut polytope is exact. This bound is tight since Laurent showed in [Lau03] that at least $\lceil n/2 \rceil$ levels are needed.

Theorem 2D. *The $\lceil n/2 \rceil$ level of the Lasserre hierarchy for the cut polytope CUT_n (as considered in [Lau03]) is exact.*

- Consider now the case $G = \mathbb{Z}_N$ and $\mathcal{S} = \{-d, -(d-1), \dots, d-1, d\}$. Here the moment polytope $\mathcal{M}(G, \mathcal{S})$ is the *trigonometric cyclic polytope* of degree d which we denote by $TC(N, 2d)$:

$$TC(N, 2d) = \text{conv}\{M(2\pi x/N) : x = 0, 1, \dots, N-1\} \subset \mathbb{R}^{2d}, \quad (6)$$

where $M(\theta)$ is the degree d trigonometric moment curve:

$$M(\theta) = (\cos(\theta), \sin(\theta), \cos(2\theta), \sin(2\theta), \dots, \cos(d\theta), \sin(d\theta)).$$

When interpreted from the dual point of view, Theorem 3 shows that $TC(N, 2d)$ has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$.

Theorem 3D. *Let N and d be two integers and assume that d divides N . The trigonometric cyclic polytope $TC(N, 2d)$ defined in (6) has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$.*

Note that in the case $d = 1$ the polytope $TC(N, 2d)$ is nothing but the regular N -gon in the plane. Theorem 3D thus recovers, and extends to the case where N is not a power of two, a result from [FSP14] giving a semidefinite lift of the regular N -gon of size $O(\log N)$.

For $d > 1$ our result is, as far as we are aware, the first nontrivial semidefinite programming lift of a cyclic polytope. Furthermore, in the regime where $N = d^2$ our lift is provably smaller than any linear programming lift: Indeed, since $TC(d^2, 2d)$ is d -neighborly [Gal63], a lower bound from [FKPT13] concerning neighborly polytopes shows that any linear programming lift of $TC(d^2, 2d)$ must have size at least $\Omega(d^2)$, whereas our semidefinite programming lift in this case has size $O(d \log d) = o(d^2)$. To the best of our knowledge this gives the first example of a family of polytopes $(P_d)_{d \in \mathbb{N}}$ in increasing dimensions where $\text{xc}_{\text{PSD}}(P_d) = o(\text{xc}_{\text{LP}}(P_d))$ where xc_{PSD} and xc_{LP} are respectively the SDP and LP extension complexity (see Section 2.3 for the definitions). More precisely, we have:

Corollary 1. *There exists a family $(P_d)_{d \in \mathbb{N}}$ of polytopes where $P_d \subset \mathbb{R}^{2d}$ such that*

$$\frac{\text{xc}_{\text{PSD}}(P_d)}{\text{xc}_{\text{LP}}(P_d)} = O\left(\frac{\log d}{d}\right).$$

The only nontrivial linear programming lift for cyclic polytopes that we are aware of is a construction by Bogomolov et al. [BFMP14] for the polytope $\text{conv}\{(i, i^2, \dots, i^d) : i = 1, \dots, N\}$ which has size $(\log N)^{\lfloor d/2 \rfloor}$.

Main ideas We now briefly describe the main ideas behind Theorem 1, which can be summarized in three steps:

1. *A sum-of-squares certificate with a sparse Gram matrix:* Given a nonnegative function $f : G \rightarrow \mathbb{R}_+$ it is easy to see, since G is finite, that f can be written as a sum-of-squares. When the function f is supported on \mathcal{S} , one can show that f admits a specific sum-of-squares representation where the *Gram matrix* Q , in the basis of characters, is sparse according to the graph $\text{Cay}(\widehat{G}, \mathcal{S})$.
2. *Chordal completion:* Let Γ be a chordal cover of the graph $\text{Cay}(\widehat{G}, \mathcal{S})$. Using well-known results concerning positive semidefinite matrices that are sparse according to a chordal graph [GT84, GJSW84] (see Section 2.2 for more details) one can decompose the Gram matrix Q into a sum of positive-semidefinite matrices, where each matrix is supported on a maximal clique of Γ . In terms of sum-of-squares representation, this means that the function f can be written as:

$$f = \sum_j |f_j|^2 \quad (7)$$

where each f_j is supported on a maximal clique \mathcal{C}_j of Γ .

3. *Translation of cliques:* The problem with the decomposition (7) is that even though each maximal clique \mathcal{C}_j might be small, the union of the \mathcal{C}_j 's might be large, and thus the total support of (7) might be large (in fact the union of the \mathcal{C}_j is the whole \widehat{G}). In order to reduce the total support of the sum-of-squares certificate (7), we use the following simple but crucial observation: if h is a function supported on \mathcal{C} and if $\chi \in \widehat{G}$ then χh is supported on $\chi\mathcal{C}$ and we have $|\chi h|^2 = |h|^2$. Thus if for each maximal clique \mathcal{C}_j of Γ we choose a certain $\chi_j \in \widehat{G}$ then, by translating each term in (7) by χ_j we obtain a sum-of-squares representation of f of the form $f = \sum_j |\tilde{h}_j|^2$ where \tilde{h}_j is supported in $\chi_j\mathcal{C}_j$. Having chosen the χ_j such that $\chi_j\mathcal{C}_j \subseteq \mathcal{T}$ for all maximal cliques \mathcal{C}_j (cf. Theorem 1), we get a representation of f as a sum-of-squares of functions supported on \mathcal{T} .

Organization The paper is organized as follows. Section 2 starts by giving a brief review of Fourier analysis of finite abelian groups, as well as a review of chordal graphs, chordal covers and the main results concerning decomposition/matrix completion with chordal sparsity structure [GT84, GJSW84]. In Section 3 we prove our main result, Theorem 1. We present the proof using the two dual viewpoints of sum-of-squares certificates and in terms of moment polytopes. In Section 4 we look at the case of the hypercube $G = \{-1, 1\}^n$ mentioned earlier, and we look in particular at quadratic functions on the hypercube. We give an explicit chordal cover for the corresponding Cayley graph and we show how it leads to a proof of Laurent's conjecture. In Section 5 we look at the special case $G = \mathbb{Z}_N$ and functions of degree d . We give an explicit chordal cover for the corresponding graphs, and we discuss the consequences concerning positive semidefinite lifts of the trigonometric cyclic polytope.

Notations We collect some of the notations used in the paper. If $z \in \mathbb{C}$ we denote by \bar{z} the complex conjugate of z . Given a square matrix $X \in \mathbb{C}^{n \times n}$ the Hermitian conjugate of X is denoted X^* , and X is called Hermitian if $X^* = X$. The space of $n \times n$ Hermitian matrices is denoted \mathbf{H}^n and the cone of Hermitian positive semidefinite matrices is denoted by \mathbf{H}_+^n . Similarly we denote by \mathbf{S}^n the space of $n \times n$ real symmetric matrices and by \mathbf{S}_+^n the cone of $n \times n$ real symmetric positive semidefinite matrices. If V is an arbitrary set, we will denote by \mathbb{C}^V the space of complex vectors indexed by elements of V , and by \mathbf{H}^V the space of Hermitian matrices where rows and columns are indexed by elements of V (and similarly for \mathbf{H}_+^V and $\mathbf{S}^V, \mathbf{S}_+^V$).

2 Preliminaries

In this section we present some background material needed for the paper: we first recall some of the basic results and terminology concerning Fourier analysis on finite abelian groups [Rud90, Ter99], then we review the definition of chordal graph and the main results concerning sparse positive semidefinite matrices and matrix completion. We also review some of the terminology concerning lifts of polytopes/extended formulations.

2.1 Fourier analysis on finite groups

Let G be a finite abelian group which we denote multiplicatively, and let $\mathcal{F}(G, \mathbb{C})$ be the vector space of complex-valued functions on G . A *character* χ of G is a *group homomorphism* $\chi : G \rightarrow (\mathbb{C}^*, \times)$, i.e., it is an element of $\mathcal{F}(G, \mathbb{C})$ which satisfies:

$$\chi(xy) = \chi(x)\chi(y) \quad \forall x, y \in G.$$

Since G is abelian, one can easily show that the (pointwise) product of two characters is a character and that the (pointwise) inverse of a character is again a character. Thus if we denote by \widehat{G} the set of characters of G , then \widehat{G} forms an abelian group, where the group operation corresponds to pointwise multiplication. The group \widehat{G} is known as the *dual group* of G . Observe that since G is finite, if χ is a character then for any $x \in G$ we have $\chi(x)^{|G|} = \chi(x^{|G|}) = \chi(1_G) = 1$, which implies that $|\chi(x)| = 1$. It follows that the inverse of a character χ is simply its (pointwise) complex conjugate $\bar{\chi}$.

Consider the standard inner product on $\mathcal{F}(G, \mathbb{C})$:

$$\langle f, g \rangle = \frac{1}{|G|} \sum_{x \in G} \overline{f(x)} g(x) \quad \forall f, g \in \mathcal{F}(G, \mathbb{C}). \quad (8)$$

A crucial property of the set of characters \widehat{G} is that they form an orthonormal basis of $\mathcal{F}(G, \mathbb{C})$, which is called the *Fourier basis* of G . Note that this implies in particular that $|\widehat{G}| = |G|$. We summarize this in the following theorem:

Theorem 4. *Let G be a finite abelian group and let \widehat{G} be the set of characters of G . Then \widehat{G} is an abelian group with pointwise multiplication. Furthermore $|\widehat{G}| = |G|$ and \widehat{G} forms an orthonormal basis of $\mathcal{F}(G, \mathbb{C})$ for the standard inner product (8).*

We now illustrate the previous theorem in the two examples $G = \{-1, 1\}^n$ (the hypercube) and $G = \mathbb{Z}_N$ presented in the introduction.

Example 1 (Fourier analysis on the hypercube). Let $G = \{-1, 1\}^n$ be the hypercube in dimension n which forms a group of size 2^n under componentwise multiplication, isomorphic to \mathbb{Z}_2^n . Observe that if S is a subset of $[n]$ then the function χ_S defined by:

$$\chi_S : \{-1, 1\}^n \rightarrow \mathbb{C}^*, \quad \chi_S(x) = \prod_{i \in S} x_i$$

satisfies $\chi_S(xy) = \chi_S(x)\chi_S(y)$, and thus is a character of G . For example χ_\emptyset is the constant function equal to 1, and $\chi_{[n]}$ is the function $\chi_{[n]}(x) = x_1 \dots x_n$. One can show that these are all the characters of G , i.e., $\widehat{G} = \{\chi_S, S \subseteq [n]\}$. Thus the decomposition of a function $f : \{-1, 1\}^n \rightarrow \mathbb{C}$ in the basis of characters takes the form:

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i,$$

where $\widehat{f}(S)$ are the Fourier coefficients of f . ◇

Example 2 (Fourier analysis on \mathbb{Z}_N). Let N be an integer and consider the (additive) group $G = \mathbb{Z}_N$ of integers modulo N . For $k \in \mathbb{Z}_N$, define χ_k by

$$\chi_k : \mathbb{Z}_N \rightarrow \mathbb{C}^*, \quad \chi_k(x) = e^{2i\pi kx/N}.$$

Note that χ_k satisfies $\chi_k(x+y) = \chi_k(x)\chi_k(y)$ and thus χ_k is a character of \mathbb{Z}_N . It is not hard to show that any character χ of \mathbb{Z}_N actually must have the form $\chi = \chi_k$ for some $k \in \mathbb{Z}_N$. Thus the dual group $\widehat{\mathbb{Z}_N}$ of \mathbb{Z}_N is $\widehat{\mathbb{Z}_N} = \{\chi_k, k \in \mathbb{Z}_N\}$. Note that $\chi_k \chi_{k'} = \chi_{k+k'}$ and $(\chi_k)^{-1} = \overline{\chi_k} = \chi_{-k}$, and thus $\widehat{\mathbb{Z}_N}$ is isomorphic to \mathbb{Z}_N . According to Theorem 4, any function $f : \mathbb{Z}_N \rightarrow \mathbb{C}$ can be decomposed in the basis of characters:

$$f(x) = \sum_{k \in \mathbb{Z}_N} \widehat{f}(k) e^{2i\pi kx/N} \quad \forall x \in \mathbb{Z}_N.$$

This decomposition is nothing but the well-known Fourier decomposition of discrete signals of length N . ◇

For a general finite abelian group G , the Fourier decomposition of a function $f : G \rightarrow \mathbb{C}$, in the orthonormal basis of characters takes the form:

$$f(x) = \sum_{\chi \in \widehat{G}} \widehat{f}(\chi) \chi(x).$$

The coefficients $\widehat{f}(\chi)$ are the *Fourier coefficients* of f . By orthonormality of the basis of characters, we have for any $\chi \in \widehat{G}$:

$$\widehat{f}(\chi) = \langle \chi, f \rangle = \frac{1}{|G|} \sum_{x \in G} \overline{\chi(x)} f(x).$$

The *support* of a function f , denoted $\text{supp } f$ is the set of characters χ for which $\widehat{f}(\chi) \neq 0$:

$$\text{supp } f = \{\chi \in \widehat{G} : \widehat{f}(\chi) \neq 0\}.$$

2.2 Chordal graphs and matrix completion

In this section we recall some of the main results concerning sparse matrix decomposition and matrix completion with a chordal sparsity structure. For more details, we refer the reader to [GJSW84, GT84] and [AHMR88].

Chordal graphs Let $\mathcal{G} = (V, E)$ be a graph. The graph \mathcal{G} is called *chordal* if any cycle of length at least four has a chord. A *chordal cover* (also called *triangulation*) of \mathcal{G} is a graph $\mathcal{G}' = (V, E')$ where $E \subset E'$ and where \mathcal{G}' is chordal. Figure 1 shows a non-chordal graph \mathcal{G} on four vertices and a chordal cover \mathcal{G}' of \mathcal{G} .




Figure 1. A non-chordal graph \mathcal{G} and a chordal cover \mathcal{G}' of \mathcal{G} .

A subset $\mathcal{C} \subseteq V$ is a *clique* in \mathcal{G} if $\{i, j\} \in E$ for all $i, j \in \mathcal{C}$, $i \neq j$. The clique \mathcal{C} is called *maximal* if it is not a strict subset of another clique \mathcal{C}' of \mathcal{G} . For example the maximal cliques of the graph \mathcal{G}' shown in Figure 1 are $\{1, 2, 4\}$ and $\{2, 3, 4\}$.

Sparse matrices Let $Q \in \mathbf{H}^V$ be a Hermitian positive semidefinite matrix where rows and columns are indexed by some set V . Assume furthermore that Q is sparse according to some graph $\mathcal{G} = (V, E)$, i.e.,

$$Q_{ij} \neq 0, i \neq j \Rightarrow \{i, j\} \in E.$$

One of the main tools used in this paper is a result from [GT84, GJSW84] which allows to decompose sparse positive semidefinite matrices as a sum of positive semidefinite matrices supported on a small subset of rows/columns. We say that a Hermitian matrix A is supported on $\mathcal{C} \subseteq V$ if $A_{ij} = 0$ whenever $i \notin \mathcal{C}$ or $j \notin \mathcal{C}$. The result can be stated as follows:

Theorem 5. ([GT84, GJSW84]) *Let Q be a Hermitian positive semidefinite matrix, and assume that Q is sparse according to some graph \mathcal{G} . Assume furthermore that \mathcal{G} is chordal. Then for every maximal clique \mathcal{C} of \mathcal{G} there exists a Hermitian positive semidefinite matrix $Q_{\mathcal{C}}$ supported on \mathcal{C} such that:*

$$Q = \sum_{\mathcal{C}} Q_{\mathcal{C}}. \quad (9)$$

Remark. If the sparsity pattern \mathcal{G} of Q is not chordal, one can still apply the previous theorem by considering a chordal cover \mathcal{G}' of \mathcal{G} . Indeed if Q is sparse according to \mathcal{G} then it also clearly sparse according to \mathcal{G}' , since $\mathcal{G} \subseteq \mathcal{G}'$. In this case the summation (9) is over the maximal cliques of \mathcal{G}' .

Example 3. We can illustrate the previous theorem with a simple 4×4 matrix. Let Q be the 4×4 Hermitian positive semidefinite matrix given by:

$$Q = \begin{bmatrix} 2 & 1-i & 0 & 1+i \\ 1+i & 2 & 1-i & 0 \\ 0 & 1+i & 2 & 1-i \\ 1-i & 0 & 1+i & 2 \end{bmatrix}.$$

Note that Q is sparse according to the “square graph” \mathcal{G} shown in Figure 1(left). Since \mathcal{G} is not chordal we cannot directly apply Theorem 5 with \mathcal{G} , but we can apply it with \mathcal{G}' shown in Figure 1(right) which is a chordal cover of \mathcal{G} . In this case Theorem 5 asserts that one can decompose Q as a sum of two positive

semidefinite matrices supported respectively on the maximal cliques, $\{1, 2, 4\}$ and $\{2, 3, 4\}$. For this example, it is not hard to find an explicit decomposition, for example we can verify that:

$$Q = \underbrace{\begin{bmatrix} 2 & 1-i & 0 & 1+i \\ 1+i & 1 & 0 & i \\ 0 & 0 & 0 & 0 \\ 1-i & -i & 0 & 1 \end{bmatrix}}_{\succeq 0} + \underbrace{\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 1-i & -i \\ 0 & 1+i & 2 & 1-i \\ 0 & i & 1+i & 1 \end{bmatrix}}_{\succeq 0}.$$

◇

Matrix completion One can also state Theorem 5 in its dual form, in terms of the *matrix completion problem*. Given a graph $\mathcal{G} = (V, E)$, a \mathcal{G} -*partial matrix* X is a matrix where only the diagonal entries, as well as the entries X_{ij} for $\{i, j\} \in E$ are specified. Given a \mathcal{G} -partial matrix X , the positive semidefinite matrix completion problem asks whether X can be completed into a full $|V| \times |V|$ Hermitian matrix that is positive semidefinite. Clearly a necessary condition for such a completion to exist is that $X[\mathcal{C}, \mathcal{C}] \succeq 0$ for all cliques \mathcal{C} of \mathcal{G} (note that if \mathcal{C} is a clique of \mathcal{G} , then all the entries of $X[\mathcal{C}, \mathcal{C}]$ are specified). When \mathcal{G} is chordal, it turns out that this condition is also sufficient. The following theorem can actually be obtained from Theorem 5 via duality:

Theorem 6. ([GJSW84]) *Let $\mathcal{G} = (V, E)$ be a graph and let X be a \mathcal{G} -partial matrix. Assume that \mathcal{G} is chordal. Then X can be completed into a full $|V| \times |V|$ Hermitian positive semidefinite matrix if, and only if, $X[\mathcal{C}, \mathcal{C}] \succeq 0$ for all maximal cliques \mathcal{C} of \mathcal{G} .*

2.3 Lifts of polytopes

In this section we recall some of the definitions and terminology concerning *lifts (or extended formulations) of polytopes*. The concepts defined here are not used in the proofs of our theorems, but simply make some of the results more convenient to state. We refer the reader to [Yan91, GPT13] for more details.

Let $P \subset \mathbb{R}^d$ be a polytope. We say that P has a *LP lift* of size k if P can be expressed as the linear projection of an affine section of the cone \mathbb{R}_+^k , i.e., if there exist $\pi : \mathbb{R}^k \rightarrow \mathbb{R}^d$ linear and an affine subspace $L \subset \mathbb{R}^k$ such that:

$$P = \pi(\mathbb{R}_+^k \cap L). \quad (10)$$

Note that this definition is equivalent to say that P is the projection of a polytope Q with k facets. The smallest k such that P has a LP lift of size k is called the *LP extension complexity* of P and is denoted $\text{xc}_{\text{LP}}(P)$.

The definition of LP lift can be extended to *PSD lifts*, where instead we are looking to describe P using *linear matrix inequalities*. Formally, we say that P has a *Hermitian PSD lift* of size k if P can be expressed as the linear projection of an affine section of the Hermitian positive semidefinite cone \mathbf{H}_+^k , i.e., if there exist $\pi : \mathbf{H}^k \rightarrow \mathbb{R}^d$ linear, and an affine subspace $L \subset \mathbf{H}^k$ such that:

$$P = \pi(\mathbf{H}_+^k \cap L). \quad (11)$$

The smallest k for which P has a PSD lift of size k is called the *PSD extension complexity* of P and denoted $\text{xc}_{\text{PSD}}(P)$. Note that one can also define PSD lifts with the cone of real symmetric positive semidefinite matrices \mathbf{S}_+^k (instead of \mathbf{H}_+^k) and in this case we call the lift a *real PSD lift*.

3 Main result for general finite abelian groups

In this section we state and prove our main result in the general setting of finite abelian groups G . We first describe the primal point of view concerning sparse sum-of-squares certificates of nonnegative functions, and then we present the dual point of view related to moment polytopes.

3.1 Nonnegative functions and sum-of-squares certificates

Let G be a finite abelian group and let $\mathcal{F}(G, \mathbb{C})$ be the space of complex-valued functions on G . Given a nonnegative function $f : G \rightarrow \mathbb{R}_+$, a *sum-of-squares certificate* for f takes the form:

$$f(x) = \sum_{k=1}^K |f_k(x)|^2 \quad \forall x \in G. \quad (12)$$

where $f_1, \dots, f_K \in \mathcal{F}(G, \mathbb{C})$.

It is well-known in the literature on polynomial optimization (see e.g., [Par00, Nes00, Las01]) that the existence of sum-of-squares certificates can be expressed in terms of the existence of a certain positive semidefinite matrix called a *Gram* matrix for f . This connection between sum-of-squares certificates and positive semidefinite matrices will be important in this paper, and so we recall this connection more formally in the next proposition:

Proposition 1. *Let $n = |G|$ and let b_1, \dots, b_n be a basis for $\mathcal{F}(G, \mathbb{C})$. Let $f : G \rightarrow \mathbb{R}$ be a real-valued function on G . Then f has a sum-of-squares representation (12), if, and only if, there exists a $n \times n$ Hermitian positive semidefinite matrix Q such that*

$$f(x) = [b(x)]^* Q [b(x)] = \sum_{1 \leq i, j \leq n} Q_{ij} \overline{b_i(x)} b_j(x) \quad \forall x \in G \quad (13)$$

where $[b(x)] := [b_i(x)]_{i=1, \dots, n} \in \mathbb{C}^n$. If (13) holds where Q is Hermitian positive semidefinite, we say that Q is a *Gram matrix* for f in the basis b_1, \dots, b_n .

Proof. Assume first that f is a sum of squares, i.e., $f(x) = \sum_{k=1}^K |f_k(x)|^2$. Since (b_1, \dots, b_n) forms a basis of $\mathcal{F}(G, \mathbb{C})$ we can write $f_k(x) = \sum_{i=1}^n \overline{a_{ki}} b_i(x)$ for some coefficients $a_{ki} \in \mathbb{C}$. Note that $|f_k(x)|^2 = \sum_{1 \leq i, j \leq n} a_{ki} \overline{a_{kj}} \overline{b_i(x)} b_j(x)$ and thus $f(x) = \sum_k |f_k(x)|^2 = \sum_{1 \leq i, j \leq n} Q_{ij} \overline{b_i(x)} b_j(x)$ where Q is the Hermitian matrix defined by: $Q_{i,j} = \sum_k a_{ki} \overline{a_{kj}}$. Note that Q is positive semidefinite since it has the form $Q = \sum_k a_k a_k^*$ where a_k is the vector $(a_k)_i = a_{ki}$.

We now show the converse. Assume f can be written as (13). Since Q is positive semidefinite, we can find vectors a_k such that $Q = \sum_{k=1}^K a_k a_k^*$. If we define f_k to be the function $f_k(x) = \sum_{i=1}^n \overline{a_{ki}} b_i(x)$ then we can verify that $f = \sum_{k=1}^K |f_k|^2$. \square

Given $y \in G$ define the Dirac function δ_y at y by:

$$\delta_y(x) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{else.} \end{cases}$$

Then it is easy to see that we have:

Proposition 2. *Any nonnegative function f on G has a sum-of-squares certificate as:*

$$f(x) = \sum_{y \in G} |\sqrt{f(y)} \delta_y(x)|^2 \quad \forall x \in G. \quad (14)$$

Said differently, a nonnegative function f is a sum-of-squares because if we pick b_1, \dots, b_n to be the basis of Dirac functions, then f satisfies Equation (13) where Q is the diagonal matrix consisting of the values taken by f on G .

Since we are working with functions on a finite abelian group G , it is more natural (and more beneficial, as we see later) to look at sum-of-squares representation in the basis of *characters*. One reason for this is that typically the functions f we are interested in have a small support in the basis of characters and in this case one can hope to find a sum-of-squares decomposition which also only involves a small number of characters. The next theorem is simply a change-of-basis in the formula (14).

Proposition 3. Let $f : G \rightarrow \mathbb{R}$ and assume that f is nonnegative, i.e., $f(x) \geq 0$ for all $x \in G$. Define the Hermitian matrix $Q \in \mathbb{R}^{|G| \times |G|}$ indexed by characters $\chi \in \widehat{G}$ by:

$$Q_{\chi, \chi'} = \widehat{f}(\overline{\chi}\chi'). \quad (15)$$

Then Q is positive semidefinite and we have for any $x \in G$:

$$f(x) = \frac{1}{|G|} [\chi(x)]^* Q [\chi(x)] = \frac{1}{|G|} \sum_{\chi, \chi' \in \widehat{G}} Q_{\chi, \chi'} \overline{\chi(x)} \chi'(x). \quad (16)$$

Proof. Consider the matrix $X = [\chi(x)]_{x \in G, \chi \in \widehat{G}}$ where rows are indexed by elements $x \in G$ and columns are indexed by characters $\chi \in \widehat{G}$. Since the characters form an orthonormal basis of $\mathcal{F}(G, \mathbb{C})$ for the inner product (8), this means that the matrix $\frac{1}{\sqrt{|G|}} X$ is a unitary matrix. Note that we can rewrite the definition (15) of Q in matrix terms as follows:

$$Q = \frac{1}{|G|} X^* \text{diag}([f(x)]_{x \in G}) X,$$

where $\text{diag}([f(x)]_{x \in G})$ is the diagonal matrix with the values $f(x)$ on the diagonal. This shows that the eigenvalues of Q are the values $\{f(x), x \in G\}$, and thus Q is positive semidefinite. Since $\frac{1}{\sqrt{|G|}} X$ is unitary we also get that:

$$\text{diag}([f(x)]_{x \in G}) = \frac{1}{|G|} X Q X^*.$$

which, when evaluated at the diagonal entries is exactly Equation (16). \square

Example 4. We now include a simple example to illustrate the previous theorem. Let $G = \mathbb{Z}_6$ and consider the function

$$f(x) = 1 - \frac{1}{2}(\chi_1(x) + \chi_{-1}(x)) = 1 - \cos(2\pi x/6) \quad \forall x \in \mathbb{Z}_6. \quad (17)$$

Clearly $f(x) \geq 0$ for all $x \in \mathbb{Z}_6$. Also note that $\widehat{f}(0) = 1$, $\widehat{f}(1) = \widehat{f}(-1) = -1/2$ and $\widehat{f}(k) = 0$ for all $k \notin \{-1, 0, 1\}$. The matrix Q defined in (15) associated to this function f takes the form:

$$Q = \begin{bmatrix} 1 & -1/2 & 0 & 0 & 0 & -1/2 \\ -1/2 & 1 & -1/2 & 0 & 0 & 0 \\ 0 & -1/2 & 1 & -1/2 & 0 & 0 \\ 0 & 0 & -1/2 & 1 & -1/2 & 0 \\ 0 & 0 & 0 & -1/2 & 1 & -1/2 \\ -1/2 & 0 & 0 & 0 & -1/2 & 1 \end{bmatrix} \quad (18)$$

\diamond

We are now interested in nonnegative functions $f : G \rightarrow \mathbb{R}$ that are supported on a subset $\mathcal{S} \subseteq \widehat{G}$, i.e., $\widehat{f}(\chi) = 0$ for all $\chi \notin \mathcal{S}$. For such functions we are interested in finding *sparse* sum-of-squares certificates for f , i.e., we are interested in finding a set $\mathcal{T} \subseteq \widehat{G}$ such that any nonnegative function f supported on \mathcal{S} has a sum-of-squares certificate of the form:

$$f = \sum_{k=1}^K |f_k|^2 \quad \text{where} \quad \text{supp } f_k \subseteq \mathcal{T} \quad \forall k = 1, \dots, K. \quad (19)$$

The main idea to obtain such a “sparse” sum-of-squares certificate of f is to exploit the sparsity of the Gram matrix Q from Proposition 3. Indeed, note that if $\text{supp } f = \mathcal{S}$, then the Gram matrix Q of Proposition 3 has a specific sparsity structure:

$$Q_{\chi, \chi'} \neq 0 \Leftrightarrow \overline{\chi}\chi' \in \mathcal{S}.$$

In other words, the sparsity structure of Q is given by the *Cayley graph* $\text{Cay}(\widehat{G}, \mathcal{S})$. Recall the definition of a Cayley graph:

Definition 1. Let H be a group and let $\mathcal{S} \subset H$ be a subset of H that is symmetric, i.e., $x \in \mathcal{S} \Rightarrow x^{-1} \in \mathcal{S}$. The Cayley graph $\text{Cay}(H, \mathcal{S})$ is the graph where vertices are the elements of the group H , and where two distinct vertices $x, y \in H$ are connected by an edge if $x^{-1}y \in \mathcal{S}$ (or $y^{-1}x \in \mathcal{S}$, which is the same since \mathcal{S} is symmetric).

Remark. We do not require the set \mathcal{S} to be a generator for the group H and hence the graph $\text{Cay}(H, \mathcal{S})$ may be disconnected. Also observe that the set $\mathcal{S} = \text{supp } f$ in our case is symmetric since f is real-valued; indeed when f is real-valued we have $\widehat{f}(\bar{\chi}) = \overline{\widehat{f}(\chi)}$ for all $\chi \in \widehat{G}$ and thus $\chi \in \text{supp } f \Rightarrow \bar{\chi} \in \text{supp } f$.

To obtain a set $\mathcal{T} \subseteq \widehat{G}$ such that (19) holds for all functions f supported on \mathcal{S} we will study *chordal covers* of the graph $\text{Cay}(\widehat{G}, \mathcal{S})$. We now introduce the key definition of *Fourier support* for a graph with vertices \widehat{G} .

Definition 2. Let Γ be a graph with vertices \widehat{G} . We say that Γ has *Fourier support* (or *frequencies*) $\mathcal{T} \subseteq \widehat{G}$ if for any maximal clique \mathcal{C} of Γ there exists $\chi_{\mathcal{C}} \in \widehat{G}$ such that $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$ (where $\chi_{\mathcal{C}}\mathcal{C} := \{\chi_{\mathcal{C}}\chi : \chi \in \mathcal{C}\}$ is the translation of \mathcal{C} by $\chi_{\mathcal{C}}$).

Note that one can also state the definition of Fourier support of Γ in terms of *equivalence classes* of cliques: Given a subset $\mathcal{C} \subseteq \widehat{G}$ define the *equivalence class* of \mathcal{C} to be all the subsets of \widehat{G} that can be obtained from \mathcal{C} by translation, i.e., it is the set $[\mathcal{C}] := \{\chi\mathcal{C} : \chi \in \widehat{G}\}$. Using this terminology, the graph Γ has Fourier support \mathcal{T} if for any maximal clique \mathcal{C} of Γ there is at least one representative from $[\mathcal{C}]$ that is contained in \mathcal{T} .

We are now ready to state and prove our main theorem (the theorem below was stated as Theorem 1 in the introduction and we reuse the same numbering here since it is just a restatement).

Theorem 1. Let \mathcal{S} be a symmetric subset of \widehat{G} and assume that $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support $\mathcal{T} \subseteq \widehat{G}$. Then any nonnegative function supported on \mathcal{S} admits a sum-of-squares certificate supported on \mathcal{T} .

Proof. Let $f : G \rightarrow \mathbb{R}$ be a nonnegative function supported on \mathcal{S} . Let Q be the Gram matrix (15) associated to the sum-of-squares representation of f in the basis of characters. We saw that Q is sparse according to the Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$. Since Γ is a cover of $\text{Cay}(\widehat{G}, \mathcal{S})$, Q is also sparse according to Γ . Thus, since Γ is chordal, using Theorem 5 we can find a decomposition of Q as follows:

$$Q = \sum_{\mathcal{C}} Q_{\mathcal{C}} \quad (20)$$

where the sum is over the maximal cliques \mathcal{C} of Γ and where each $Q_{\mathcal{C}}$ is a positive semidefinite matrix supported on \mathcal{C} . Note that Equation (20) implies that for all $x \in G$:

$$[\chi(x)]^* Q [\chi(x)] = \sum_{\mathcal{C}} [\chi(x)]^* Q_{\mathcal{C}} [\chi(x)],$$

where $[\chi(x)] := [\chi(x)]_{\chi \in \widehat{G}}$. Since $f(x) = [\chi(x)]^* Q [\chi(x)] / |G|$ the above equation says that:

$$f(x) = \sum_{\mathcal{C}} f_{\mathcal{C}}(x)$$

where we let $f_{\mathcal{C}}(x) := [\chi(x)]^* Q_{\mathcal{C}} [\chi(x)] / |G|$. Since $Q_{\mathcal{C}}$ is positive semidefinite and supported on \mathcal{C} , this means that each $f_{\mathcal{C}}(x)$ is a sum-of-squares of functions supported on $\mathcal{C} \subseteq \widehat{G}$, i.e.,

$$f_{\mathcal{C}} = \sum_k |f_{\mathcal{C},k}|^2$$

where $\text{supp } f_{\mathcal{C},k} \subseteq \mathcal{C}$.

According to Definition 2, we know that there exist $\chi_{\mathcal{C}} \in \widehat{G}$ for each maximal clique \mathcal{C} of Γ such that $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$. Now, observe that:

$$f = \sum_{\mathcal{C}} f_{\mathcal{C}} = \sum_{\mathcal{C}} \sum_k |f_{\mathcal{C},k}|^2 \stackrel{(i)}{=} \sum_{\mathcal{C}} \sum_k |\chi_{\mathcal{C}} f_{\mathcal{C},k}|^2 \stackrel{(ii)}{=} \sum_{\mathcal{C}} \sum_k |\widetilde{f}_{\mathcal{C},k}|^2$$

where in (i) we used the fact that $|\chi_C|^2 = 1$ and in (ii) we let $\tilde{f}_{C,k} = \chi_C f_{C,k}$ which is supported on $\chi_C \mathcal{C} \subseteq \mathcal{T}$. Thus we have shown that f is a sum-of-squares of functions supported on \mathcal{T} . \square

Example 5. Let $G = \mathbb{Z}_6$ and let $\mathcal{S} = \{-1, 0, 1\} \subset \widehat{\mathbb{Z}_6}$. We will use the previous theorem to show that any nonnegative function on \mathbb{Z}_6 supported on $\mathcal{S} = \{-1, 0, 1\}$ is a sum-of-squares of functions supported on $\mathcal{T} = \{-1, 0, 1, 3\} \subseteq \widehat{\mathbb{Z}_6}$. The Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ is the cycle graph on 6 nodes shown in Figure 2(left). Clearly the graph is not chordal since the cycle $0, 1, \dots, 5$ has no chord. Figure 2(right) shows a chordal cover Γ of $\text{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ where the maximal cliques are:

$$\mathcal{C}_1 = \{0, 1, 3\}, \quad \mathcal{C}_2 = \{1, 2, 3\}, \quad \mathcal{C}_3 = \{3, 4, 5\}, \quad \mathcal{C}_4 = \{0, 3, 5\}.$$




Figure 2. Left: The Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ is the cycle graph on 6 nodes. Right: A chordal cover of the cycle graph, Γ .

Observe that if we translate the clique $\mathcal{C}_2 = \{1, 2, 3\}$ by -2 we get $\{-1, 0, 1\}$ and similarly if we translate the clique $\{3, 4, 5\}$ by -4 we also get $\{-1, 0, 1\}$. Thus by choosing

$$\chi_{\mathcal{C}_1} = 0, \quad \chi_{\mathcal{C}_2} = -2, \quad \chi_{\mathcal{C}_3} = -4, \quad \chi_{\mathcal{C}_4} = 0$$

we get that $\chi_C + \mathcal{C} \subseteq \{-1, 0, 1, 3\}$ for all maximal cliques \mathcal{C} of Γ (we used the fact that $5 = -1$ in \mathbb{Z}_6). In other words we have shown that Γ is a chordal cover of $\text{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 0, 1\})$ with frequencies $\{-1, 0, 1, 3\}$. Thus by Theorem 1, this means that any nonnegative function on \mathbb{Z}_6 supported on $\{-1, 0, 1\}$ can be written as a sum-of-squares of functions supported on $\{-1, 0, 1, 3\}$. \diamond

3.2 Dual point of view: moment matrices and matrix completion

Section 3.1 shows that nonnegative functions $f : G \rightarrow \mathbb{R}$ that have Fourier support \mathcal{S} can be written as sums of Hermitian squares of functions $h : G \rightarrow \mathbb{C}$ with Fourier support \mathcal{T} , where \mathcal{T} satisfies a combinatorial property related to the Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$. In this section we describe the same results from the dual point of view, arriving at a dual statement (Theorem 1D) of the main result of Section 3.1 (Theorem 1). The dual result describes a certain moment polytope $\mathcal{M}(G, \mathcal{S})$ (see Definition 3 to follow) as the projection of a section of the cone of positive semidefinite matrices indexed by \mathcal{T} . We could obtain this dual result by applying a conic duality argument directly to the statement of Theorem 1. The purpose of this section, however, is to re-explain the results of Section 3.1 from an alternative viewpoint.

3.2.1 Moment polytopes

We begin by describing the moment polytope $\mathcal{M}(G, \mathcal{S})$ where $\mathcal{S} \subseteq \widehat{G}$ is a collection of characters. Concretely $\mathcal{M}(G, \mathcal{S})$ is the convex hull of a collection of complex vectors indexed by \mathcal{S} :

$$\mathcal{M}(G, \mathcal{S}) = \text{conv}\{(\chi(x))_{\chi \in \mathcal{S}} \in \mathbb{C}^{\mathcal{S}} : x \in G\}.$$

The following definition is equivalent, somewhat easier to work with, and more readily generalizable.

Definition 3. The *moment polytope* $\mathcal{M}(G, \mathcal{S})$ with respect to the characters $\mathcal{S} \subseteq \widehat{G}$, is the convex polytope

$$\mathcal{M}(G, \mathcal{S}) = \{(\mathbb{E}_\mu[\chi])_{\chi \in \mathcal{S}} \in \mathbb{C}^{\mathcal{S}} : \mu \text{ a probability measure supported on } G\}.$$

For example, in the case where $G = \mathbb{Z}_6$ and $\mathcal{S} = \{-1, 1\}$, the moment polytope is

$$\mathcal{M}(\mathbb{Z}_6, \{-1, 1\}) = \text{conv}\{(e^{\frac{-2\pi ki}{6}}, e^{\frac{2\pi ki}{6}}) : k \in \mathbb{Z}_6\}.$$

Nonnegative functions on finite abelian groups are sums of squares. Furthermore, we can express sums of squares concretely in terms of a Gram matrix. There is a similarly concrete way to describe the constraints that must be satisfied by a collection of complex numbers $\ell \in \mathbb{C}^{\widehat{G}}$ if they are a valid collection of moments of a probability measure supported on G . This description is given naturally in terms of a matrix constructed from $(\ell_\chi)_{\chi \in \widehat{G}}$.

Definition 4. If $\ell \in \mathbb{C}^{\widehat{G}}$, the associated *moment matrix* is the square matrix with rows and columns indexed by \widehat{G} of the form

$$[M(\ell)]_{\chi, \chi'} = \ell_{\overline{\chi}\chi'} \quad \text{for all } \chi, \chi' \in \widehat{G}.$$

If $\mathcal{T} \subseteq \widehat{G}$ and $\ell \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}}$, the associated *truncated moment matrix* is the square matrix with rows and columns indexed by \mathcal{T} of the form

$$[M_{\mathcal{T}}(\ell)]_{\chi, \chi'} = \ell_{\overline{\chi}\chi'} \quad \text{for all } \chi, \chi' \in \mathcal{T}.$$

We now describe the dual version of the fact that any nonnegative function on G is a sum of squares. Writing this in coordinates gives a concrete description in terms of moment matrices. That probability measures are real-valued and nonnegative, and have total mass one corresponds to the conditions that for any valid moment vector ℓ , the moment matrix $M(\ell)$ is Hermitian and positive semidefinite, and has unit diagonal.

Proposition 4. *The moment polytope can be expressed as*

$$\mathcal{M}(G, \widehat{G}) = \{(\ell(\chi))_{\chi \in \widehat{G}} : \ell(1_{\widehat{G}}) = 1, \ell(|f|^2) \geq 0 \text{ for all } f \in \mathcal{F}(G, \mathbb{C})\}.$$

Equivalently, defining coordinates $\ell_\chi := \ell(\chi)$ with respect to the character basis for $\mathcal{F}(G, \mathbb{C})$ we have

$$\mathcal{M}(G, \widehat{G}) = \{\ell \in \mathbb{C}^{\widehat{G}} : \ell_{1_{\widehat{G}}} = 1, M(\ell) \succeq 0\}.$$

Proof. See Appendix A.1 □

Some readers may recognize this as essentially a statement of Bochner's theorem for finite abelian groups [Rud90]. If we regard ℓ as a function $\ell : \widehat{G} \rightarrow \mathbb{C}$, the condition that $M(\ell) \succeq 0$ is exactly saying that ℓ is a *positive definite function* on \widehat{G} . While we could use the language of positive definite functions throughout this section, we instead use the more concrete language of moment matrices. We do this both so that our descriptions are compatible with the literature on polynomial optimization, and so that they are easy to implement in code.

The polytope $\mathcal{M}(G, \mathcal{S})$ is just the projection of $\mathcal{M}(G, \widehat{G})$ onto the coordinates corresponding to \mathcal{S} . Alternatively we can think of $\mathcal{M}(G, \mathcal{S})$ as those points in $\mathbb{C}^{\mathcal{S}}$ that can be completed to valid moment sequences. This suggests describing $\mathcal{M}(G, \mathcal{S})$ in terms of a structured positive semidefinite matrix completion problem. In this problem the diagonal is given, the entries corresponding to the edges of $\text{Cay}(\widehat{G}, \mathcal{S})$ are given, and the goal is to complete the matrix to a positive semidefinite moment matrix.

Corollary 2. *The moment polytope with respect to \mathcal{S} can be expressed as*

$$\mathcal{M}(G, \mathcal{S}) = \{\ell \in \mathbb{C}^{\mathcal{S}} : \exists y \in \mathbb{C}^{\widehat{G}} \text{ s.t. } \ell_\chi = y_\chi \text{ for all } \chi \in \mathcal{S}, y_{1_{\widehat{G}}} = 1, M(y) \succeq 0\}. \quad (21)$$

In the example of $\mathcal{M}(\mathbb{Z}_6, \{-1, 1\})$ we index characters by $\{0, 1, 2, 3, 4, 5\}$ (so that $-1 = 5$). Then Corollary 2 tells us that $\mathcal{M}(\mathbb{Z}_6, \{-1, 1\})$ is the set of (ℓ_5, ℓ_1) such that

$$\exists y_2, y_3, y_4 \in \mathbb{C} \text{ s.t. } \begin{bmatrix} 1 & \ell_1 & y_2 & y_3 & y_4 & \ell_5 \\ \ell_5 & 1 & \ell_1 & y_2 & y_3 & y_4 \\ y_4 & \ell_5 & 1 & \ell_1 & y_2 & y_3 \\ y_3 & y_4 & \ell_5 & 1 & \ell_1 & y_2 \\ y_2 & y_3 & y_4 & \ell_5 & 1 & \ell_1 \\ \ell_1 & y_2 & y_3 & y_4 & \ell_5 & 1 \end{bmatrix} \succeq 0.$$

Note that we adopt the convention, throughout, that writing $M(y) \succeq 0$ implies that $M(y)$ is Hermitian. This may, in turn, imply certain linear equalities on y . For instance, in the above example the notation implies that $\ell_1 = \overline{\ell_5} = \overline{\ell_{-1}}$ and $y_2 = \overline{y_4} = \overline{y_{-2}}$ and $y_3 = \overline{y_3}$.

3.2.2 Moment polytopes and matrix completion

In Proposition 3 we saw that any nonnegative (and hence sum of squares) function with Fourier support \mathcal{S} has a Gram matrix that is sparse with respect to the graph $\text{Cay}(\widehat{G}, \mathcal{S})$. The dual statement is that the moment polytope $\mathcal{M}(G, \mathcal{S})$ can be described in terms of an *unstructured* positive semidefinite matrix completion problem. In this case we are given the diagonal and the entries corresponding to the edges of $\text{Cay}(\widehat{G}, \mathcal{S})$ and just need to complete the matrix to be positive semidefinite, without requiring it to be a moment matrix.

Proposition 5. *The moment polytope with respect to \mathcal{S} can be expressed as*

$$\mathcal{M}(G, \mathcal{S}) = \{\ell \in \mathbb{C}^{\mathcal{S}} : \exists Y \in \mathbf{H}_+^{\widehat{G}} \text{ s.t. } Y_{\chi, \chi'} = \ell_{\overline{\chi}\chi'} \text{ whenever } \overline{\chi}\chi' \in \mathcal{S} \text{ and } Y_{\chi, \chi} = 1 \text{ for all } \chi \in \widehat{G}\}.$$

Idea of proof. We describe the main idea of the proof, giving the details in Appendix A.1. The key issue is to show that if $\ell \in \mathbb{C}^{\mathcal{S}}$ has a completion to a positive semidefinite matrix $Y \in \mathbf{H}_+^{\widehat{G}}$ then ℓ also has a completion to a positive semidefinite *moment* matrix $M(y)$ (for some $y \in \mathbb{C}^{\widehat{G}}$). This can be established by observing that the group \widehat{G} acts on the rows and columns of Hermitian matrices $\mathbf{H}^{\widehat{G}}$ by permutations. This action fixes, pointwise, positive semidefinite moment matrices. Averaging the orbit of Y under this group action gives a moment matrix with the desired properties. \square

Recall that Theorem 1 gave a combinatorial condition under which any nonnegative function with Fourier support \mathcal{S} is not just a sum of squares, but a sum of squares of functions with Fourier support \mathcal{T} . The dual version says that under the same combinatorial condition, to check that $\ell \in \mathcal{M}(G, \mathcal{S})$ we are not required to complete it to a full positive semidefinite moment matrix. Instead it is enough to complete it to a truncated moment matrix $M_{\mathcal{T}}(y)$ for some $y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}}$.

Theorem 1D. *Let \mathcal{S} be a symmetric subset of \widehat{G} . If $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support $\mathcal{T} \subseteq \widehat{G}$ then*

$$\mathcal{M}(G, \mathcal{S}) = \{\ell \in \mathbb{C}^{\mathcal{S}} : \exists y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}} \text{ s.t. } y_{\chi} = \ell_{\chi} \text{ for all } \chi \in \mathcal{S}, y_{1_{\widehat{G}}} = 1, \text{ and } M_{\mathcal{T}}(y) \succeq 0\}. \quad (22)$$

Idea of proof. Again we summarize the key idea, deferring a detailed proof to Appendix A.1. The main issue is to show that being able to complete ℓ to a truncated moment matrix $M_{\mathcal{T}}(y)$ for some $y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}}$ implies we can complete it to a positive semidefinite matrix $Y \in \mathbf{H}_+^{\widehat{G}}$. Proposition 5 then implies we can complete it to a positive semidefinite moment matrix.

The key observation, analogous to the translation of frequencies idea in the proof of Theorem 1, is that given $y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}}$, the partial moment matrices $M_{\chi\mathcal{T}}(y)$ for any $\chi \in \widehat{G}$ are all the same. Hence imposing that $M_{\mathcal{T}}(y) \succeq 0$, implies that $M_{\chi\mathcal{T}}(y) \succeq 0$ for all $\chi \in \widehat{G}$. Since $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover Γ with Fourier support \mathcal{T} , we construct from y a partial matrix supported on the maximal cliques of Γ . The conditions that $M_{\chi\mathcal{T}}(y) \succeq 0$ for all $\chi \in \widehat{G}$ are enough to show that all the principal submatrices supported on maximal cliques of Γ are positive semidefinite. The chordal matrix completion result (Theorem 6) completes the proof. \square

Example (Example 5 cont.). Recall that the Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_6}, \{-1, 1\})$ is the 6-cycle. We label the elements of $\widehat{\mathbb{Z}_6}$ by $\{0, 1, 2, 3, 4, 5\}$. The Cayley graph has a chordal cover (see Figure 2) with Fourier support $\mathcal{T} = \{-1, 0, 1, 3\} = \{5, 0, 1, 3\}$ (in \mathbb{Z}_6). Observe that $\mathcal{T}^{-1}\mathcal{T} = \widehat{\mathbb{Z}_6}$. Applying Theorem 1D we see that

$$\mathcal{M}(\mathbb{Z}_6, \{-1, 1\}) = \left\{ (\ell_5, \ell_1) : \exists y_2, y_3, y_4 \in \mathbb{C} \text{ s.t. } \begin{bmatrix} 1 & \ell_1 & y_3 & \ell_5 \\ \ell_5 & 1 & y_2 & y_4 \\ y_3 & y_4 & 1 & y_2 \\ \ell_1 & y_2 & y_4 & 1 \end{bmatrix} \succeq 0 \right\}.$$

3.3 Real sums-of-squares and moment polytopes

The main results of Sections 3.1 and 3.2 work with sums of Hermitian squares of complex-valued functions and complex Hermitian moment matrices respectively. While this is convenient mathematically, computationally it is desirable to work with real-valued functions and real symmetric moment matrices. In this section we give real versions of Theorem 1 and Theorem 1D. These are the forms most suited to implementation and the forms we use when discussing the examples in Sections 4 and 5 to follow.

Basic observations The main additional observation we make is that the dual group \widehat{G} consists of two types of characters: those that are real-valued, and those that are not. It is helpful to think of this decomposition in terms of the involution $\chi \mapsto \chi^{-1} = \overline{\chi}$ on \widehat{G} . Real-valued characters are those that are fixed by inversion (complex conjugation). The remaining characters come in inverse (complex conjugate) pairs.

To fix notation let \widehat{G}_0 denote the real-valued characters and $\widehat{G}_{-1} \cup \widehat{G}_1$ be a fixed partition of the remaining characters into conjugate pairs. In particular $\widehat{G}_{-1} = \widehat{G}_1^{-1}$. If $\mathcal{S} \subseteq \widehat{G}$ is symmetric (i.e. $\mathcal{S}^{-1} = \mathcal{S}$) then inversion restricts to an involution on \mathcal{S} and so we have the decomposition $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{S}_{-1} \cup \mathcal{S}_1$ where $\mathcal{S}_i = \widehat{G}_i \cap \mathcal{S}$ for $i = -1, 0, 1$.

Sums-of-squares Suppose f is a sum of Hermitian squares of functions $f_j : G \rightarrow \mathbb{C}$, each supported on $\mathcal{T} \subseteq \widehat{G}$. Then f can be expressed as a sum-of-squares of real-valued functions $\text{Re}[f_j]$ and $\text{Im}[f_j]$ as

$$f = \sum_j |f_j|^2 = \sum_j (\text{Re}[f_j]^2 + \text{Im}[f_j]^2).$$

Clearly $\text{Re}[f_j]$ and $\text{Im}[f_j]$ are supported on the symmetric subset $\mathcal{T} \cup \mathcal{T}^{-1}$ of characters. As such, the real analogue of Theorem 1 is the following, the only modification being that we insist that \mathcal{T} is symmetric.

Theorem 7. *Let $\mathcal{S} \subseteq \widehat{G}$ be symmetric and let $f : G \rightarrow \mathbb{R}$ be a non-negative function supported on \mathcal{S} . If $\mathcal{T} \subseteq \widehat{G}$ is symmetric and $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support \mathcal{T} then f is a sum of squares of real-valued functions supported on \mathcal{T} .*

Real moment matrices and moment polytopes We now develop the real analogue of moment polytopes and moment matrices. The discussion is more explicit (and more involved) than was required for the sum-of-squares viewpoint. We begin by defining real moment polytopes. These are just linear transformations of the moment polytopes of Section 3.2. Throughout this section for any symmetric $\mathcal{S} \subseteq \widehat{G}$ partitioned as $\mathcal{S}_0 \cup \mathcal{S}_{-1} \cup \mathcal{S}_1$ fix a linear map $\mathcal{R}_{\mathcal{S}} : \mathbb{C}^{\mathcal{S}} \rightarrow \mathbb{R}^{|\mathcal{S}|}$ defined by

$$\mathcal{R}_{\mathcal{S}}(\ell) = ((\ell_{\chi})_{\chi \in \mathcal{S}_0}, (\text{Re}[\ell_{\chi}], \text{Im}[\ell_{\chi}])_{\chi \in \mathcal{S}_1}).$$

Observe that $\mathcal{R}_{\mathcal{S}}$ depends on the partitioning of \mathcal{S} .

Definition 5. If $\mathcal{S} \subseteq \widehat{G}$ is symmetric and partitioned as $\mathcal{S}_0 \cup \mathcal{S}_{-1} \cup \mathcal{S}_1$, the *real moment polytope* with respect to \mathcal{S} is the image of $\mathcal{M}(G, \mathcal{S})$ under $\mathcal{R}_{\mathcal{S}}$, i.e.

$$\mathcal{M}^{\mathbb{R}}(G, \mathcal{S}) = \mathcal{R}_{\mathcal{S}}(\mathcal{M}(G, \mathcal{S})).$$

This is a polytope in $\mathbb{R}^{|\mathcal{S}|}$ that is affinely isomorphic to $\mathcal{M}(G, \mathcal{S})$. For example, in the case where $G = \mathbb{Z}_6$ and $\mathcal{S} = \{-1, 1\}$ is decomposed as $\mathcal{S}_{-1} = \{-1\}$ and $\mathcal{S}_1 = \{1\}$, the real moment polytope is

$$\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_6, \{-1, 1\}) = \text{conv}\{(\cos(\frac{2\pi k}{6}), \sin(\frac{2\pi k}{6})) : k \in \mathbb{Z}_6\},$$

the regular hexagon in the plane.

We have a description of $\mathcal{M}(G, \mathcal{S})$, and hence of $\mathcal{M}^{\mathbb{R}}(G, \mathcal{S})$ in terms of Hermitian positive semidefinite matrices. Using this it is straightforward to give a description in terms of real symmetric positive semidefinite matrices of twice the size (via (31) in Appendix A.2). Our aim is to describe the real moment polytope $\mathcal{M}^{\mathbb{R}}(G, \mathcal{S})$ in terms of real symmetric positive semidefinite matrices *without increasing the size of the description*. It turns out that we can do this whenever \mathcal{T} has a property that is related to, but less restrictive than, being symmetric.

Definition 6. A subset $\mathcal{T} \subseteq \widehat{G}$ has an *equalizing involution* if there is an involution $\sigma : \mathcal{T} \rightarrow \mathcal{T}$ such that $\sigma(\chi)\chi = \sigma(\chi')\chi'$ for all $\chi, \chi' \in \widehat{G}$.

Observe that if \mathcal{T} is symmetric, then the map that sends every element of \mathcal{T} to its inverse satisfies $\chi^{-1}\chi = (\chi')^{-1}\chi' = 1_{\widehat{G}}$ for all $\chi, \chi' \in \mathcal{T}$. Hence any symmetric subset of \widehat{G} has an equalizing involution. An example of a set that has an equalizing involution but is not symmetric is $\mathcal{T} = \{0, 1, 2, 3\} \subseteq \widehat{\mathbb{Z}_5}$. Then \mathcal{T} is not symmetric and, furthermore, there is no $k \in \widehat{\mathbb{Z}_5}$ such that $k + \mathcal{T}$ is symmetric. Nevertheless, \mathcal{T} does have an equalizing involution given by $\sigma(0) = 3$, $\sigma(3) = 0$, $\sigma(1) = 2$ and $\sigma(2) = 1$ since $\sigma(k) + k = 3$ for all $k \in \mathcal{T}$.

Our main result in this section is the following.

Theorem 8. Let $\mathcal{S} \subseteq \widehat{G}$ be symmetric and let $\mathcal{T} \subseteq \widehat{G}$ have an equalizing involution σ . If $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support \mathcal{T} then

$$\begin{aligned} \mathcal{M}^{\mathbb{R}}(G, \mathcal{S}) = \{ \mathcal{R}_{\mathcal{S}}(\ell) \in \mathbb{R}^{|\mathcal{S}|} : \exists y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}} \text{ s.t. } y_{\bar{\chi}} = \overline{y_{\chi}} \text{ for all } \chi \in \mathcal{T}^{-1}\mathcal{T}, \\ y_{\chi} = \ell_{\chi} \text{ for all } \chi \in \mathcal{S}, y_{1_{\widehat{G}}} = 1, \text{ and } [\text{Re}[y_{\bar{\chi}\chi'}] - \text{Im}[y_{\sigma(\bar{\chi})\chi'}]]_{\chi, \chi' \in \mathcal{T}} \succeq 0 \}. \end{aligned}$$

Note that the main change between Theorem 8 and Theorem 1D is that we have replaced the condition that the (Hermitian) truncated moment matrix $M_{\mathcal{T}}(y)$ be positive semidefinite with the condition that a *real symmetric* matrix indexed by \mathcal{T} be positive semidefinite. We also explicitly add the conjugate symmetry constraint that $y_{\bar{\chi}} = \overline{y_{\chi}}$, which was implied by $M_{\mathcal{T}}(y)$ being Hermitian in Theorem 1D.

Example 6. Before giving a proof, we apply Theorem 8 to the case of $\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_6, \{-1, 1\})$, the regular hexagon in the plane. Recall that in this case we can take $\mathcal{T} = \{0, 1, 3, 5\}$ which is symmetric. Hence \mathcal{T} has $\sigma(\chi) = \chi^{-1}$ as an equalizing involution. Decomposing \widehat{G} as $\widehat{G}_0 = \{0, 3\}$ and, for instance $\widehat{G}_1 = \{1, 2\}$ and $\widehat{G}_{-1} = \{4, 5\}$ we have that $\mathcal{S}_0 = \{0\}$ and $\mathcal{S}_1 = \{1\}$. Note also that $\mathcal{T}^{-1}\mathcal{T} = \widehat{G}$ in this case, giving $(\mathcal{T}^{-1}\mathcal{T})_i = \widehat{G}_i$ for $i = -1, 0, 1$. Ordering the elements of \mathcal{T} as $(0, 1, 3, 5)$ and writing $u_j = \text{Re}[\ell_j]$, $v_j = \text{Im}[\ell_j]$, $w_j = \text{Re}[y_j]$ and $x_j = \text{Im}[y_j]$, we see that

$$[\text{Re}[y_{\bar{\chi}\chi'}] - \text{Im}[y_{\sigma(\bar{\chi})\chi'}]]_{\chi, \chi' \in \mathcal{T}} = [w_{\bar{\chi}\chi'} - x_{\chi\chi'}]_{\chi, \chi' \in \mathcal{T}} = \begin{bmatrix} w_0 - x_0 & w_1 - x_1 & w_3 - x_3 & w_5 - x_5 \\ w_5 - x_1 & w_0 - x_2 & w_2 - x_4 & w_4 - x_0 \\ w_3 - x_3 & w_4 - x_4 & w_0 - x_0 & w_2 - x_2 \\ w_1 - x_5 & w_2 - x_0 & w_4 - x_2 & w_0 - x_4 \end{bmatrix}.$$

The conjugate symmetry constraint on y implies that $x_0 = x_3 = 0$, $w_1 = w_5$, $x_1 = -x_5$, $w_2 = w_4$, and $x_2 = -x_4$. Applying these and the constraints that $u_1 = w_1$, $v_1 = x_1$ and $w_0 = 1$, we obtain a description of the regular hexagon in the plane as

$$\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_6, \{-1, 1\}) = \left\{ (u_1, v_1) : \exists w_2, w_3, x_2 \in \mathbb{R} \text{ s.t. } \begin{bmatrix} 1 & u_1 - v_1 & w_3 & u_1 + v_1 \\ u_1 - v_1 & 1 - x_2 & w_2 + x_2 & w_2 \\ w_3 & w_2 + x_2 & 1 & w_2 - x_2 \\ u_1 + v_1 & w_2 & w_2 - x_2 & 1 + x_2 \end{bmatrix} \succeq 0 \right\}.$$

◇

Proof of Theorem 8. We use a novel result that allows us to write certain complex Hermitian linear matrix inequalities as real symmetric linear matrix inequalities of the same size. The key requirement is that restricted to the subspace of Hermitian matrices of interest, entry-wise complex conjugation can be expressed as congruence by an orthogonal symmetric matrix.

Lemma 1. *Let \mathcal{L} be a subspace (over the reals) of \mathbf{H}^d . Suppose there is some orthogonal $J \in O(d)$ such that $J^2 = I$ and*

$$JLJ^T = \overline{L} \quad \text{for all } L \in \mathcal{L},$$

i.e. congruence by J restricted to \mathcal{L} is entry-wise complex conjugation. Then

$$\{L \in \mathcal{L} : L \in \mathbf{H}_+^d\} = \{L \in \mathcal{L} : \text{Re}[L] - J \text{Im}[L] \in \mathbf{S}_+^d\}.$$

Proof. See Appendix A.2. □

Let \mathcal{L} be the subspace (over the reals) of $\mathbf{H}^{\mathcal{T}}$ given by

$$\mathcal{L} = \{M(y) : y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}}, y_{\overline{\chi}} = \overline{y_{\chi}}, \text{ for all } \chi \in \mathcal{T}^{-1}\mathcal{T}\}.$$

Let J be the $|\mathcal{T}| \times |\mathcal{T}|$ permutation matrix representing the equalizing involution $\sigma : \mathcal{T} \rightarrow \mathcal{T}$. Since σ is an involution, J satisfies $J^2 = I$. The definition of equalizing involution comes from our desire that σ satisfy the relation $\sigma(\chi)\overline{\sigma(\chi')} = \overline{\chi}\chi'$, equivalent to the defining relation in Definition 6, for all $\chi, \chi' \in \mathcal{T}$. Then for any conjugate symmetric y ,

$$\begin{aligned} JM(y)J &= [y_{\overline{\chi}\chi'}]_{\sigma(\chi), \sigma(\chi') \in \mathcal{T}} = [y_{\overline{\sigma(\chi)\sigma(\chi')}}]_{\chi, \chi' \in \mathcal{T}} \\ &= [\overline{y}_{\sigma(\chi)\sigma(\chi')}]_{\chi, \chi' \in \mathcal{T}} \quad \text{by conjugate symmetry of } y \\ &= [\overline{y}_{\overline{\chi}\chi'}]_{\chi, \chi' \in \mathcal{T}} \quad \text{since } \sigma \text{ is an equalizing involution.} \end{aligned}$$

Hence congruence by J corresponds to entry-wise complex conjugation restricted to \mathcal{L} . Applying Lemma 1 we can conclude that for all $y \in \mathbb{C}^{\mathcal{T}^{-1}\mathcal{T}}$,

$$M(y) \succeq 0 \iff y_{\overline{\chi}} = \overline{y_{\chi}} \quad \forall \chi \in \mathcal{T}^{-1}\mathcal{T} \quad \text{and} \quad [\text{Re}[y_{\overline{\chi}\chi'}] - \text{Im}[y_{\sigma(\overline{\chi})\chi'}]]_{\chi, \chi' \in \mathcal{T}} \succeq 0.$$

This completes the proof of Theorem 8. □

4 Application 1: The case $G = \mathbb{Z}_2^n$ and binary quadratic optimization

In this section we apply the results of Section 3.1 to the case of nonnegative quadratic forms on the vertices of the hypercube in n dimensions. Dually, the moment polytope of interest in this section is the n th cut polytope CUT_n . Our main aim is to establish Laurent's conjecture [Lau03, Conjecture 4] that any nonnegative quadratic form on the vertices of the hypercube in n dimension is a sum of squares of polynomials of degree at most $\lceil n/2 \rceil$.

4.1 Quadratic forms on $\{-1, 1\}^n$ and the cut polytope

Let $G = \{-1, 1\}^n$ be the vertices of the hypercube in dimension n . View G as a group (isomorphic to \mathbb{Z}_2^n) under componentwise multiplication. Recall that the characters of G are indexed by subsets $S \in 2^{[n]}$ and are the square-free monomials

$$\chi_S(x) = \prod_{i \in S} x_i \quad \text{for all } x \in G.$$

We focus on characterizing nonnegative quadratic functions on G . These are of particular interest because the problem of maximizing a quadratic form over G i.e.

$$\max_{x \in G} \sum_{1 \leq i < j \leq n} A_{ij} x_i x_j \tag{23}$$




Figure 3. The Cayley graph $\text{Cay}(\widehat{G}, \mathcal{S})$ for $G = \{-1, 1\}^4$ and $\mathcal{S} = \{S : |S| = 0 \text{ or } |S| = 2\}$. The two connected components are $\mathcal{T}_{\text{even}}$ (left) and \mathcal{T}_{odd} (right). The vertices of \mathcal{T}_{odd} are arranged to correspond to their images in $\mathcal{T}_{\text{even}}$ under the graph automorphism $\phi(S) = \{1\} \triangle S$. We can obtain a chordal cover of $\text{Cay}(\widehat{G}, \mathcal{S})$ by forming maximal cliques on the vertices of $\mathcal{T}_{\text{even}}$ marked with filled circles, the vertices of $\mathcal{T}_{\text{even}}$ marked with open circles, and the images in \mathcal{T}_{odd} of these two cliques under the map ϕ .

includes, for example, the MAX-CUT problem, which arises when the symmetric matrix A_{ij} is the Laplacian of a (weighted) graph on n vertices. We can solve (23) by finding the smallest upper bound on the objective:

$$\min_{\gamma} \gamma \quad \text{s.t.} \quad \gamma - \sum_{1 \leq i < j \leq n} A_{ij} x_i x_j \geq 0 \quad \text{for all } x \in G. \quad (24)$$

If we have a characterization of nonnegative functions on G with Fourier support $\mathcal{S} = \{S \in 2^{[n]} : |S| = 0 \text{ or } |S| = 2\}$ as sums of squares of functions supported on $\mathcal{T} \subseteq \widehat{G}$ then we can solve (24) by solving a semidefinite optimization problem of size $|\mathcal{T}|$.

The dual picture to (24) is to consider optimization over the moment polytope $\mathcal{M}(\{-1, 1\}^n, \mathcal{S} \setminus \{\emptyset\})$, known as the *cut polytope*

$$\text{CUT}_n := \mathcal{M}(\{-1, 1\}^n, \mathcal{S} \setminus \{\emptyset\}) = \text{conv} \{(x_i x_j)_{1 \leq i < j \leq n} : (x_1, x_2, \dots, x_n) \in \{-1, 1\}^n\}.$$

We can solve the binary quadratic optimization problem (23) by optimizing the linear function defined by A over CUT_n , i.e. by solving the linear program

$$\max_{(\ell_{ij})_{1 \leq i < j \leq n}} \sum_{i < j} \ell_{ij} A_{ij} \quad \text{s.t.} \quad (\ell_{ij})_{1 \leq i < j \leq n} \in \text{CUT}_n.$$

If we have a PSD lift of the cut polytope CUT_n of size $|\mathcal{T}|$ then we can solve this optimization problem by solving a semidefinite optimization problem of size $|\mathcal{T}|$.

4.2 The associated Cayley graph

To apply the results of Section 3.1 we need to understand the graph $\text{Cay}(\widehat{G}, \mathcal{S})$. In the case $n = 4$ this graph is shown in Figure 3. Throughout this section we identify the character $\chi_S \in \widehat{G}$ with the subset $S \subseteq [n]$ that indexes it and work exclusively in the language of subsets. As such, the vertex set of $\text{Cay}(\widehat{G}, \mathcal{S})$ is $2^{[n]}$, the collection of all subsets $S \subseteq [n]$. There is an edge between two subsets S, T if and only if $|S \triangle T| = 2$. This graph is often called the *half-cube graph*.

The group operation on characters is multiplication of functions, which corresponds to taking the symmetric difference of the subsets that index the characters. In other words, if $S, T \subseteq [n]$ then

$$\chi_S(x) \chi_T(x) = \chi_{S \triangle T}(x)$$

where $S \triangle T = (S \setminus T) \cup (T \setminus S)$. As such, there is an action of \widehat{G} on the vertices of $\text{Cay}(\widehat{G}, \mathcal{S})$ by $S \cdot T = S \triangle T$. Furthermore if $\mathcal{T} \subseteq 2^{[n]}$ is a subset of the vertices of $\text{Cay}(\widehat{G}, \mathcal{S})$ we write $S \triangle \mathcal{T} := \{S \triangle T : T \in \mathcal{T}\}$.

We now record some simple observations that follow directly from the adjacency relation in $\text{Cay}(\widehat{G}, \mathcal{S})$. For convenience of notation, for $k = 0, 1, \dots, n$ let

$$\mathcal{T}_k = \{S \subseteq [n] : |S| = k\}.$$

Any edge of $\text{Cay}(\widehat{G}, \mathcal{S})$ either has both endpoints in \mathcal{T}_k for some k or one endpoint in \mathcal{T}_k and the other in \mathcal{T}_{k+2} for some k . Consequently, $\text{Cay}(\widehat{G}, \mathcal{S})$ has two connected components

$$\mathcal{T}_{\text{even}} = \mathcal{T}_0 \cup \mathcal{T}_2 \cup \dots \cup \mathcal{T}_{2\lfloor n/2 \rfloor} \quad \text{and} \quad \mathcal{T}_{\text{odd}} = \mathcal{T}_1 \cup \mathcal{T}_3 \cup \dots \cup \mathcal{T}_{2\lfloor n/2 \rfloor - 1}.$$

Define a map $\phi : 2^{[n]} \rightarrow 2^{[n]}$ by $\phi(S) = \{1\} \triangle S$. Since $|\phi(S) \triangle \phi(T)| = |S \triangle T|$ for all $S, T \in 2^{[n]}$ it follows that ϕ extends to an automorphism of $\text{Cay}(\widehat{G}, \mathcal{S})$ that exchanges $\mathcal{T}_{\text{even}}$ and \mathcal{T}_{odd} .

4.3 Applying Theorem 1

To apply Theorem 1 from Section 3.1 we need to find a subset $\mathcal{T} \subseteq 2^{[n]}$ of vertices such that $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support \mathcal{T} . The following result explicitly describes such a collection of vertices.

Proposition 6. *The graph $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support*

$$\mathcal{T} = \begin{cases} \mathcal{T}_0 \cup \mathcal{T}_2 \cup \dots \cup \mathcal{T}_{\lfloor n/2 \rfloor} & \text{if } \lfloor n/2 \rfloor \text{ even} \\ \mathcal{T}_1 \cup \mathcal{T}_3 \cup \dots \cup \mathcal{T}_{\lfloor n/2 \rfloor} & \text{if } \lfloor n/2 \rfloor \text{ odd.} \end{cases} \quad (25)$$

Proof. We give a detailed proof in Appendix A.3. \square

Example 7. To give the flavor of the proof, we discuss the case $n = 4$. In this case $\text{Cay}(\widehat{G}, \mathcal{S})$ is shown in Figure 3. Define Γ to be the graph with vertex set $2^{[4]}$ and with edges between $S, T \in \mathcal{T}_{\text{even}}$ if and only if $||S| - |T|| \leq 2$, and edges between $S, T \in \mathcal{T}_{\text{odd}}$ if and only if $||\phi(S)| - |\phi(T)|| \leq 2$. The graph Γ is chordal, with maximal cliques given by $\mathcal{C}_0 = \mathcal{T}_0 \cup \mathcal{T}_2$, $\mathcal{C}_2 = \mathcal{T}_2 \cup \mathcal{T}_4$, $\phi(\mathcal{C}_0)$, and $\phi(\mathcal{C}_2)$. The vertices in cliques \mathcal{C}_0 and \mathcal{C}_2 are indicated by open and filled circles respectively in Figure 3. (The vertices in cliques $\phi(\mathcal{C}_0)$ and $\phi(\mathcal{C}_2)$ are similarly marked.) If $\mathcal{T} = \mathcal{T}_0 \cup \mathcal{T}_2$ then we can see that Γ is a chordal cover of $\text{Cay}(\widehat{G}, \mathcal{S})$ with Fourier support \mathcal{T} by observing that $\emptyset \triangle \mathcal{C}_0 \subseteq \mathcal{T}$, $\{1, 2, 3, 4\} \triangle \mathcal{C}_2 \subseteq \mathcal{T}$, $\phi(\emptyset) \triangle \phi(\mathcal{C}_0) \subseteq \mathcal{T}$ and $\phi(\{1, 2, 3, 4\}) \triangle \phi(\mathcal{C}_2) \subseteq \mathcal{T}$. \diamond

Laurent's conjecture follows directly from Proposition 6 and Theorem 1.

Theorem 2. *Suppose $f(x) = A_\emptyset + \sum_{1 \leq i < j \leq n} A_{ij} x_i x_j$ is nonnegative on G . Then there is a collection $(h_k)_{k=1}^{|\mathcal{T}|}$ of functions $h_k : G \rightarrow \mathbb{R}$ each supported on \mathcal{T} (defined in (25)) such that*

$$f(x) = \sum_{k=1}^{|\mathcal{T}|} h_k(x)^2.$$

Consequently, any nonnegative quadratic form on G is a sum of squares of functions of degree at most $\lfloor n/2 \rfloor$.

Proof. The first assertion follows directly from Proposition 6 and Theorem 1. The second assertion holds simply because every function supported on \mathcal{T} has degree at most $\lfloor n/2 \rfloor$. \square

The dual version of this result gives a PSD lift of the cut polytope of size $|\mathcal{T}|$. It follows directly from Proposition 6 and Theorem 1D, and the observation that in this case all the characters are real-valued.

Corollary 3. *The cut polytope CUT_n has a real PSD lift of size $|\mathcal{T}|$ given by*

$$\text{CUT}_n = \left\{ \ell \in \mathbb{R}^{S \setminus \emptyset} : \exists y \in \mathbb{R}^{\mathcal{T} \triangle \mathcal{T}} \text{ s.t. } y_\emptyset = 1, \ y_{\{i,j\}} = \ell_{\{i,j\}} \text{ for } 1 \leq i < j \leq n, \text{ and } [y_{S \triangle T}]_{S,T \in \mathcal{T}} \succeq 0 \right\}$$

where \mathcal{T} is defined in (25).

In the language used, for example, in [Lau03] when discussing the Lasserre hierarchy for binary quadratic optimization, Corollary 3 could be expressed simply as $Q_{\lfloor n/2 \rfloor} = \text{CUT}_n$ for all n .

5 Application 2: The case $G = \mathbb{Z}_N$ and cyclic polytopes

In this section we apply the results of Section 3 to the case where $G = \mathbb{Z}_N$ is the (additive) group of integers modulo N . As we will see, this will allow us to obtain a positive semidefinite lift of size $O(d \log(N/d))$ for the regular trigonometric cyclic polytope with N vertices of degree d , when d divides N .

Recall from Section 2, that the characters of \mathbb{Z}_N are indexed by $k \in \mathbb{Z}_N$ and are given by:

$$\chi_k(x) = e^{2i\pi kx/N} \quad \forall x \in \mathbb{Z}_N.$$

Thus the Fourier decomposition of a function $f : \mathbb{Z}_N \rightarrow \mathbb{C}$ is given by:

$$f(x) = \sum_{k \in \mathbb{Z}_N} \hat{f}(k) e^{2i\pi kx}.$$

Furthermore, we say that a function f has degree d if it is supported on $\{-d, -(d-1), \dots, d-1, d\}$.

5.1 The case $\mathcal{S} = \{-1, 0, 1\}$: the cycle graph

In this section we are interested in obtaining sparse sum-of-squares certificates for functions of degree 1 on \mathbb{Z}_N , i.e., functions supported on $\mathcal{S} = \{-1, 0, 1\}$. Note that the real moment polytope $\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_N, \{-1, 1\})$ in this case is just the regular N -gon in the plane:

$$\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_N, \{-1, 1\}) = \text{conv} \left\{ (\cos(2\pi x/N), \sin(2\pi x/N)) : x \in \mathbb{Z}_N \right\}.$$

To obtain sparse sum-of-squares for nonnegative functions of degree 1 we are going to study the Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_N}, \{-1, 0, 1\})$. Note that this is simply the cycle graph on N vertices, which we will denote by C_N for simplicity. The object of this section is to show that this graph admits a chordal cover with a small number of frequencies. We show:

Theorem 9. *Let N be a positive integer greater than 2. Then C_N has a chordal cover with frequencies $\mathcal{T} \subseteq \widehat{\mathbb{Z}_N}$ where $|\mathcal{T}| \leq 3 \log_2 N$. More precisely the set \mathcal{T} can be described explicitly as follows: Let $k_1 < k_2 < \dots < k_l$ be the positions of the nonzero digits in the binary expansion of N so that $N = \sum_{j=1}^l 2^{k_j}$. Let k be the largest integer such that $2^k < N$ (i.e., $k = k_1 - 1$ if N is a power of two and $k = k_l$ otherwise). Then the set \mathcal{T} is given by*

$$\mathcal{T} = \{0\} \cup \{-2^i, 2^i, i = 0, \dots, k-1\} \cup \left\{ \sum_{j=1}^i 2^{k_j}, i = 1, \dots, l-2 \right\}. \quad (26)$$

Proof. The chordal cover is constructed by induction on N , see Appendix A.4 for the details. Figure 4 shows the chordal cover for $N = 8$ and $N = 16$. \square

If we combine the previous theorem with Theorem 1, we get that any nonnegative degree-1 function on \mathbb{Z}_N has a sum-of-squares certificate supported on \mathcal{T} where $|\mathcal{T}| \leq 3 \log N$. Note that this corresponds to Theorem 3 from the introduction for the case $d = 1$. Dually, this allows us to obtain a Hermitian positive semidefinite lift of the regular N -gon of size $|\mathcal{T}| \leq 3 \log N$.

In a previous paper [FSP14] we showed that the $N = 2^n$ -gon admits a positive semidefinite lift of size $2n - 1$. In fact we showed that any linear function on \mathbb{Z}_N that is nonnegative can be written as a sum-of-squares of functions supported on $\{0\} \cup \{\pm 2^i, i = 0, \dots, n-2\}$. Note that this is the same set of frequencies that we get if we plug $N = 2^n$ in (26). Thus Theorem 9 generalizes the result of [FSP14] for arbitrary N .

5.2 Degree d functions: powers of cycle graph

In this section we are interested in functions of degree d on \mathbb{Z}_N where d divides N . We show how to construct a chordal cover of the associated Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_N}, \mathcal{S})$ using the chordal cover of the cycle graph constructed in the previous section. This allows us to show that any nonnegative function on \mathbb{Z}_N of degree d has a sum-of-squares certificate supported on at most $3d \log(N/d)$ frequencies.




Figure 4. Triangulation of the 8-cycle with frequencies $\mathcal{T} = \{-2, -1, 0, 1, 2\}$ and of the 16-cycle with frequencies $\mathcal{T} = \{-4, -2, -1, 0, 1, 2, 4\}$.

5.2.1 Triangulating the Cayley graph

Observe that the Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_N}, \mathcal{S})$ when $\mathcal{S} = \{-d, \dots, d\}$ is the d 'th power of the cycle graph C_N . Recall the definition of power of a graph:

Definition 7. Let $\mathcal{G} = (V, E)$ be a graph. Given $d \in \mathbb{N}$, the d 'th power of \mathcal{G} is the graph $\mathcal{G}^d = (V, E^d)$ where two vertices $u, v \in V$ are connected if there is a path of length $\leq d$ connecting u and v in \mathcal{G} .

Following this observation, we will use the symbol C_N^d to denote the Cayley graph $\text{Cay}(\widehat{\mathbb{Z}_N}, \{-d, \dots, d\})$. Figure 5(left) shows the graph C_8^2 for $N = 8$ and $d = 2$.




Figure 5. Left: The second power of the cycle graph on 8 nodes: two nodes are connected by an edge if their distance in the cycle graph is at most 2. Right: The graph $C_4 \boxtimes K_2$. Note that $C_8^2 \subset C_4 \boxtimes K_2$. The edges in $C_4 \boxtimes K_2$ that are not in C_8^2 are indicated with a heavier line.

To construct a triangulation of C_N^d we will actually use the triangulation of the cycle graph C_N constructed in the previous section. For this, we need the following definition of *strong product* of graphs:

Definition 8. Given graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$ define the *strong product* of \mathcal{G} and \mathcal{G}' , denoted $\mathcal{G} \boxtimes \mathcal{G}'$ to be the graph with vertex set $V \times V'$ and where two vertices $(u, u') \in V \times V'$ and $(v, v') \in V \times V'$ are connected if:

$$(u = v \text{ and } \{u', v'\} \in E') \quad \text{or} \quad (\{u, v\} \in E \text{ and } u' = v') \quad \text{or} \quad (\{u, v\} \in E \text{ and } \{u', v'\} \in E').$$

Remark. An important special case is when one of the graphs, say \mathcal{G}' , is a complete graph $\mathcal{G}' = K_m$. In this case two distinct vertices (u, u') and (v, v') in $\mathcal{G} \boxtimes K_m$ are connected if either $u = v$ or $\{u, v\} \in E(\mathcal{G})$.

Given two graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V, E')$ with the same vertex set V we say that \mathcal{G}' covers \mathcal{G} and we write $\mathcal{G} \subseteq \mathcal{G}'$ if $E \subseteq E'$. Our main observation to construct a chordal cover of C_N^d is the following:

Proposition 7. *Let N and d be two integers and assume that d divides N . Let C_N^d be the d 'th power of the cycle graph C_N and let $C_{N/d}$ be the cycle graph on N/d nodes. Then*

$$C_N^d \subseteq C_{N/d} \boxtimes K_d. \quad (27)$$

Proof. To show the inclusion (27) we first need to identify the vertices of C_N^d with those of $C_{N/d} \boxtimes K_d$. Note that the vertex set of C_N^d can be identified with \mathbb{Z}_N and the vertex set of $C_{N/d}$ can be identified with $\mathbb{Z}_{N/d}$. We also identify the vertices of K_d with $\{0, \dots, d-1\}$. By definition of \boxtimes , the vertices of $C_{N/d} \boxtimes K_d$ are $\mathbb{Z}_{N/d} \times \{0, \dots, d-1\}$. Consider the map:

$$\phi : \mathbb{Z}_{N/d} \times \{0, \dots, d-1\} \rightarrow \mathbb{Z}_N, \quad \phi(q, r) = qd + r. \quad (28)$$

This map is well-defined and gives a bijection between $\mathbb{Z}_{N/d} \times \{0, \dots, d-1\}$ and \mathbb{Z}_N . The map ϕ thus identifies vertices of C_N^d with those of $C_{N/d} \boxtimes K_d$.

We now show that, with this identification, inclusion (27) holds. We need to show that if $i, i' \in \mathbb{Z}_N$ are connected in C_N^d (i.e., $i - i' \in \{-d, \dots, d\}$) then necessarily (q, r) and (q', r') are connected in $C_{N/d} \boxtimes K_d$ (i.e., $q - q' \in \{-1, 0, 1\}$), where (q, r) and (q', r') are such that $i = \phi(q, r)$ and $i' = \phi(q', r')$. Consider for $q \in \mathbb{Z}_{N/d}$ the set of vertices of C_N^d given by $V_q = \{\phi(q, r) : r = 0, \dots, d-1\} \subset \mathbb{Z}_N$. Note that $(V_q)_{q \in \mathbb{Z}_{N/d}}$ forms a partition of the vertex set of C_N^d and that $|V_q| = d$ for all q (for example if $N = 8$ and $d = 2$ (Figure 5) $V_0 = \{0, 1\}, V_1 = \{2, 3\}, V_2 = \{4, 5\}, V_3 = \{6, 7\}$). It is easy to see that if i and i' are two adjacent vertices of C_N^d , then i and i' must be in the same group (i.e., $i, i' \in V_q$) or in adjacent group (i.e., $i \in V_q$ and $i' \in V_{q+1}$ or vice-versa). In other words this means that $q - q' \in \{-1, 0, 1\}$ which means that (q, r) and (q', r') are connected in $C_{N/d} \boxtimes K_d$. \square

The previous proposition gives a natural way to construct a chordal cover of C_N^d from that of $C_{N/d}$. Indeed if Γ is a chordal cover of $C_{N/d}$ then one can show that $\Gamma \boxtimes K_d$ is a chordal cover of C_N^d and one can also characterize the maximal cliques of $\Gamma \boxtimes K_d$ in terms of those of Γ . This is the object of the next proposition.

Proposition 8. *Let $\mathcal{G} = (V, E)$ be a graph and d be any integer.*

1. *If \mathcal{G}' is such that $\mathcal{G} \subseteq \mathcal{G}'$ then $\mathcal{G} \boxtimes K_d \subseteq \mathcal{G}' \boxtimes K_d$.*
2. *If \mathcal{G} is chordal then $\mathcal{G} \boxtimes K_d$ is chordal.*
3. *All the maximal cliques of $\mathcal{G} \boxtimes K_d$ have the form $\mathcal{C} \times K_d$ where \mathcal{C} is a maximal clique of \mathcal{G} .*

Proof. 1. The first point is clear from the definition of \boxtimes .

2. Let $(u_1, v_1) \dots (u_l, v_l)$ be a cycle in $\mathcal{G} \boxtimes K_d$ of length $l \geq 4$ where $(u_l, v_l) = (u_1, v_1)$. If there exists $i \in \{1, \dots, l-1\}$ such that $u_i = u_{i+1}$ then the edge $\{(u_i, v_i), (u_{i+2}, v_{i+2})\}$ is a chord of the cycle. Otherwise note that $u_1 \dots u_l$ is a cycle in \mathcal{G} of length ≥ 4 . Since \mathcal{G} is chordal there is $1 \leq i, j \leq l-1$ with $j - i \geq 2$ such that $\{u_i, u_j\} \in E$. In this case the edge $\{(u_i, v_i), (u_j, v_j)\}$ is a chord of the cycle.
3. The third property easily follows from the fact that if $\mathcal{C} = \{(u_i, v_i), i = 1, \dots, k\}$ is a clique in $\mathcal{G} \boxtimes K_d$ then $\{u_i, i = 1, \dots, k\} \subseteq V$ is a clique in \mathcal{G} . \square

We can now use the triangulation of the cycle graph constructed in the previous section to obtain a triangulation of C_N^d .

Proposition 9. *Let N and d be two integers and assume that d divides N . If $C_{N/d}$ has a triangulation with frequencies $\mathcal{T} \subseteq \mathbb{Z}_{N/d}$, then C_N^d has a triangulation with frequencies*

$$\mathcal{T}' = \{dk + r : k \in \mathcal{T}, r \in \{0, \dots, d-1\}\}. \quad (29)$$

Note that $|\mathcal{T}'| \leq d \cdot |\mathcal{T}|$.

Proof. Let Γ be a triangulation $C_{N/d}$ with frequencies \mathcal{T} . By definition, this means that for any maximal clique \mathcal{C} of $C_{N/d}$, there is $k_{\mathcal{C}} \in \mathbb{Z}_{N/d}$ such that $k_{\mathcal{C}} + \mathcal{C} \subseteq \mathcal{T}$.

By Proposition 8, we know that $\Gamma \boxtimes K_d$ is a chordal cover of C_N^d . Let \mathcal{C}' be a maximal clique of $\Gamma \boxtimes K_d$. By Proposition 8, we know that there exists \mathcal{C} maximal clique of Γ such that $\mathcal{C}' = \mathcal{C} \times K_d = \{dq + r : q \in \mathcal{C}, r \in \{0, \dots, d-1\}\}$. Define $k_{\mathcal{C}'} = dk_{\mathcal{C}} \in \mathbb{Z}_N$ and note that:

$$\begin{aligned} k_{\mathcal{C}'} + \mathcal{C}' &= \{dk_{\mathcal{C}} + dq + r : q \in \mathcal{C}, r \in \{0, \dots, d-1\}\} \\ &= \{(k_{\mathcal{C}} + q)d + r : q \in \mathcal{C}, r \in \{0, \dots, d-1\}\} \subseteq \mathcal{T}', \end{aligned}$$

where the last inclusion follows from the fact that $k_{\mathcal{C}} + q \in \mathcal{T}$ whenever $q \in \mathcal{C}$. We have thus shown that for any maximal clique \mathcal{C}' of $\Gamma \boxtimes K_d$, there is $k_{\mathcal{C}'} \in \mathbb{Z}_N$ such that $k_{\mathcal{C}'} + \mathcal{C}' \subseteq \mathcal{T}'$. Thus this shows that $\Gamma \boxtimes K_d$ is a chordal cover of C_N^d with frequencies \mathcal{T}' . \square

Combining Proposition 9 and the triangulation of the cycle graph from Theorem 9 we get the following corollary:

Corollary 4. *Let N and d be two integers and assume that d divides N . Then the graph C_N^d has a triangulation with frequencies $\mathcal{T} \subset \widehat{\mathbb{Z}_N}$ where $|\mathcal{T}| \leq 3d \log(N/d)$.*

Using Theorem 1, this proves Theorem 3 from the introduction concerning nonnegative functions on \mathbb{Z}_N of degree d .

Theorem 3. *Let N and d be two integers and assume that d divides N . Then there exists $\mathcal{T} \subseteq \mathbb{Z}_N$ with $|\mathcal{T}| \leq 3d \log(N/d)$ such that any nonnegative function on \mathbb{Z}_N of degree at most d has a sum-of-squares certificate supported on \mathcal{T} .*

Figure 6 shows the triangulation of C_{16}^2 obtained by triangulating C_8 using Theorem 9 and applying the strong graph product with K_2 .




Figure 6. Triangulation of the graph C_{16}^2 obtained as the strong graph product of Γ and K_2 , where Γ is the triangulation of C_8 obtained from Theorem 9 and illustrated in Figure 4(left).

5.2.2 Cyclic polytopes

Observe that the real moment polytope for $G = \mathbb{Z}_N$ and $\mathcal{S} = \{-d, \dots, d\}$ is given by:

$$\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_N, \{-d, \dots, d\}) = \text{conv} \left\{ (\cos(2\pi x/N), \sin(2\pi x/N), \dots, \cos(2\pi dx/N), \sin(2\pi dx/N)), x \in \mathbb{Z}_N \right\} \subset \mathbb{R}^{2d}.$$

This is just the regular trigonometric cyclic polytope which we abbreviate by $TC(N, 2d)$:

$$TC(N, 2d) = \text{conv} \left\{ M(2\pi x/N) : x = 0, 1, \dots, N-1 \right\} \subset \mathbb{R}^{2d}, \quad (30)$$

where $M(\theta)$ is the degree d trigonometric moment curve:

$$M(\theta) = (\cos(\theta), \sin(\theta), \cos(2\theta), \sin(2\theta), \dots, \cos(d\theta), \sin(d\theta)).$$

Cyclic polytopes play an important role in polyhedral combinatorics [Zie95] and satisfy many interesting properties. For example the celebrated *Upper Bound Theorem*, states that for any $2d$ dimensional polytope P with N vertices, $f_i(P) \leq f_i(TC(N, 2d))$ for any $i = 0, \dots, 2d$, where $f_i(P)$ is the number of i -dimensional faces of a polytope P [Zie95]. Another important property of cyclic polytopes is that they are *neighborly* [Gal63] (recall that a $2d$ -dimensional polytope P is called neighborly if any collection of d vertices of P span a face of P).

The results from this section allow us to obtain a positive semidefinite lift of $TC(N, 2d)$ of size $O(d \log(N/d))$ when d divides N . More precisely, if we combine Corollary 4 and Theorem 1D we get that $TC(N, 2d)$ has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$, proving Theorem 3D from the introduction:

Theorem 3D. *Let N and d be two integers and assume that d divides N . Then the trigonometric cyclic polytope $TC(N, 2d)$ has a Hermitian positive semidefinite lift of size at most $3d \log(N/d)$.*

Real positive semidefinite lifts Using the results of Section 3.3 one can convert the Hermitian positive semidefinite lift of $TC(N, 2d)$ into a real positive semidefinite lift of size at most $4d \log(N/d)$. Indeed, first note that in the case $d = 1$, if \mathcal{T} is the set of frequencies (26) for the cycle graph C_N then $\mathcal{T} \cup (-\mathcal{T})$ has cardinality at most $4 \log(N)$ and the set $\mathcal{T} \cup (-\mathcal{T})$ clearly has an equalizing involution since it is symmetric. Thus this shows that in the case $d = 1$, the moment polytope $\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_N, \{-1, 1\})$ has a *real* positive semidefinite lift of size at most $4 \log(N)$. For the case where $d > 1$ and d divides N it is not difficult to show that if \mathcal{T} has an equalizing involution σ then \mathcal{T}' defined in (29) also has an equalizing involution. Indeed given $k \in \mathcal{T}$ and $r \in \{0, \dots, d-1\}$, define $\sigma' : \mathcal{T}' \rightarrow \mathcal{T}'$ by $\sigma'(dk + r) = d\sigma(k) + d - r - 1$ (such a map is well-defined). Then $\sigma'(dk + r) + dk + r = d(\sigma(k) + k) + d$ which is a constant independent of k and r , and thus σ' is an equalizing involution for \mathcal{T}' . Thus using the symmetric set of frequencies for the cycle graph $C_{N/d}$ (of size $4 \log(N/d)$) we get that \mathcal{T}' has size at most $4d \log(N/d)$ and has an equalizing involution. Thus this shows that $\mathcal{M}^{\mathbb{R}}(\mathbb{Z}_N, \{-d, \dots, d\})$ has a *real* positive semidefinite lift of size at most $4d \log(N/d)$.

Comparison with LP lifts One can show that in the regime $N = \Theta(d^2)$ our positive semidefinite lift for $TC(N, 2d)$ is provably smaller than any linear programming lift of $TC(N, 2d)$. Indeed, the following lower bound on the LP extension complexity of k -neighborly polytopes was proved in [FKPT13]:

Proposition 10. ([FKPT13, Proposition 5.16]) *If P be a k -neighborly polytope with N vertices then $\text{xc}_{LP}(P) \geq \min(N, (k+1)(k+2)/2)$.*

Since $TC(N, 2d)$ is d -neighborly, if we choose for example $N = d^2$ then the previous proposition asserts that $\text{xc}_{LP}(TC(d^2, 2d)) \geq \Omega(d^2)$ whereas in this case our positive semidefinite has size $O(d \log d)$. This allows us to prove the following result giving a gap between SDP extension complexity and LP extension complexity.

Corollary 1. *There exists a family $(P_d)_{d \in \mathbb{N}}$ of polytopes where $P_d \subset \mathbb{R}^{2d}$ such that*

$$\frac{\text{xc}_{PSD}(P_d)}{\text{xc}_{LP}(P_d)} = O\left(\frac{\log d}{d}\right).$$

The only nontrivial LP lift for cyclic polytopes that we are aware of is a recent construction by Bogomolov et al. [BFMP14] for the cyclic polytope

$$C(N, d) = \text{conv}\{(i, i^2, \dots, i^d) : i = 1, \dots, N\}$$

of size $(\log N)^{\lfloor d/2 \rfloor}$. Note that this lift has smaller size than the “trivial” vertex lift of $C(N, d)$ only when $d < O((\log N)/(\log \log N))$. Their construction for $d = 2$ is based on the *reflection relations* framework of Kaibel and Pashkovich [KP11] and the case of general d is then obtained via tensor products.

6 Conclusion

In this paper we studied nonnegative functions f defined on a finite abelian group G . We looked at functions f that have a *sparse* Fourier support \mathcal{S} and we identified a certain combinatorial condition involving chordal covers of the Cayley graph $\text{Cay}(G, \mathcal{S})$, that guarantees the existence of a *sparse* sum-of-squares certificate for any nonnegative function supported on \mathcal{S} . We applied our general framework to two special cases. First we looked at quadratic functions defined on the hypercube $G = \{-1, 1\}^n$ and we showed that any nonnegative quadratic function on G has a sum-of-squares certificate of degree at most $\lceil n/2 \rceil$. This proves a conjecture by Laurent from 2003 [Lau03] and shows that the Lasserre hierarchy for the cut polytope converges after $\lceil n/2 \rceil$ steps. Second, we looked at nonnegative functions defined on $G = \mathbb{Z}_N$, the group of integers modulo N . We showed that when d divides N , any degree d nonnegative function on \mathbb{Z}_N has a sum-of-squares certificate with functions supported on some \mathcal{T} where $|\mathcal{T}| \leq O(d \log(N/d))$. From the geometric point of view, this establishes that the *regular trigonometric cyclic polytope* of degree d in N vertices has a positive semidefinite lift of size $O(d \log(N/d))$. For the regime $N = \Theta(d^2)$ this gives us a family of polytopes in increasing dimensions, where the ratio of the PSD extension complexity to the LP extension complexity is $O((\log d)/d)$.

Equivariance of lifts Observe that the moment polytopes $\mathcal{M}(G, \mathcal{S}) \subset \mathbb{C}^{\mathcal{S}}$ considered in this paper are invariant under the natural action of G on $\mathbb{C}^{\mathcal{S}}$ given by $\rho(x) = \text{diag}([\chi(x)]_{\chi \in \mathcal{S}})$. One can show that the Hermitian PSD lifts for $\mathcal{M}(G, \mathcal{S})$ considered in this paper respect this symmetry of $\mathcal{M}(G, \mathcal{S})$, in a certain formal sense known as *equivariant lifts* defined in [FSP13].

Open problems We conclude by briefly discussing two concrete problems about regular trigonometric cyclic polytopes arising from this work. In previous work we showed that any \mathbb{Z}_N -equivariant PSD lift of the regular N -gon must have size at least $\ln(N/2)$ [FSP14, Theorem 9]. Since all the Hermitian PSD lifts in this paper are equivariant, this shows that it is not possible to construct substantially better equivariant PSD lifts of the regular N -gons than the construction given in Theorem 3D with $d = 1$. It would be interesting to establish corresponding lower bounds for general d .

Problem 1. Find lower bounds on the size of (\mathbb{Z}_N -equivariant) PSD lifts of the regular trigonometric cyclic polytope of degree d in N vertices.

Problem 9.9 of [FGP⁺14] asks for a “family of polytopes that exhibits a large (e.g. exponential) gap between its nonnegative and psd ranks”. We have found that the regular trigonometric cyclic polytope of degree d in $N = d^2$ vertices gives an explicit family of polytopes for which we can prove a significant (but far from exponential) gap between PSD and LP extension complexity. It may be possible to prove a better lower bound (parameterized by N and d) on the LP extension complexity for this family of polytopes since the bound we use only uses the fact that the polytope is neighborly. Such a lower bound may allow us to establish a larger gap between the LP and PSD extension complexity of these polytopes.

Problem 2. Find $N(d)$ to make the gap between the LP and PSD extension complexity of regular trigonometric cyclic polytopes of degree d in $N(d)$ vertices as large as possible.

A Additional proofs

A.1 Proofs from Section 3.2

In this section we provide detailed proofs of the main results in Section 3.2.

Proof of Lemma 4. Suppose μ is a probability measure supported on G . Then because μ is a probability measure, $(\mathbb{E}_\mu[\chi])_{\chi \in \mathcal{S}}$ certainly satisfies $\mathbb{E}_\mu[1_{\hat{G}}] = 1$ and $\mathbb{E}_\mu[|f|^2] \geq 0$ whenever $f \in \mathcal{F}(G, \mathbb{C})$. Conversely, suppose ℓ is a linear functional on $\mathcal{F}(G, \mathbb{C})$ such that $\ell(1_{\hat{G}}) = 1$ and $\ell(|f|^2) \geq 0$ for all $f \in \mathcal{F}(G, \mathbb{C})$. We show that $\ell(\cdot)$ coincides with $\mathbb{E}_\mu[\cdot]$ for some probability measure μ supported on G . For any $x \in G$ define

$\mu(\{x\}) = \ell(\delta_x)$. Since $\delta_x = |\delta_x|^2$ we have that $\mu(\{x\}) = \ell(|\delta_x|^2) \geq 0$ for all $x \in G$. Since ℓ is linear

$$\sum_{x \in G} \mu(\{x\}) = \ell\left(\sum_{x \in G} \delta_x\right) = \ell(1_{\widehat{G}}) = 1.$$

Hence μ defines a probability measure supported on G and $\ell(\cdot)$ is exactly the corresponding expectation $\mathbb{E}_\mu[\cdot]$. The second description of $\mathcal{M}(G, \widehat{G})$ in the lemma follows by rewriting the condition $\ell(|f|^2)$ for all $f \in \mathcal{F}(G, \mathbb{C})$ in coordinates with respect to the character basis as

$$\ell\left(\left|\sum_{\chi \in \widehat{G}} \widehat{f}(\chi)\chi\right|^2\right) = \sum_{\chi, \chi' \in \widehat{G}} \ell(\overline{\widehat{f}(\chi)}\widehat{f}(\chi')) \quad \text{for all } (\widehat{f}(\chi))_{\chi \in \widehat{G}} \in \mathbb{C}^{\widehat{G}}.$$

This is exactly the definition of the matrix $M(\ell) = [\ell_{\overline{\chi\chi'}}]_{\chi, \chi' \in \widehat{G}}$ being Hermitian positive semidefinite. \square

Proof of Proposition 5. If $(\ell_\chi)_{\chi \in \mathcal{S}} \in \mathcal{M}(G, \mathcal{S})$ then from Corollary 2 there is $(y_\chi)_{\chi \in \widehat{G}}$ such that $y_\chi = \ell_\chi$ for all $\chi \in \mathcal{S}$, $y_{1_{\widehat{G}}} = 1$, and $M(y) \succeq 0$. Hence we can take $Y_{\chi, \chi'} = y_{\overline{\chi\chi'}}$ to show that $(\ell_\chi)_{\chi \in \mathcal{S}}$ is an element of the right hand side of (21).

Conversely, suppose there exists $Y \in \mathbf{H}_+^{\widehat{G}}$ with $Y_{\chi, \chi'} = \ell_{\overline{\chi\chi'}}$ whenever $\overline{\chi\chi'} \in \mathcal{S}$ and $Y_{\chi, \chi} = 1$ for all $\chi \in \widehat{G}$. Our task is to construct, from Y some $(y_\chi)_{\chi \in \widehat{G}}$ such that $y_\chi = \ell_\chi$ for all $\chi \in \mathcal{S}$, $y_{1_{\widehat{G}}} = 1$, and $M(y) \succeq 0$. Observe that \widehat{G} acts on Hermitian matrices $\mathbf{H}^{\widehat{G}}$ indexed by elements of \widehat{G} by simultaneously permuting the rows and columns, i.e. by $[\lambda \cdot Y]_{\chi, \chi'} = Y_{\overline{\lambda\chi}, \overline{\lambda\chi'}}$. We construct a new matrix Z by averaging Y over this group action:

$$Z_{\chi, \chi'} := \frac{1}{|\widehat{G}|} \sum_{\lambda \in \widehat{G}} [\lambda \cdot Y]_{\chi, \chi'} = \frac{1}{|\widehat{G}|} \sum_{\lambda \in \widehat{G}} Y_{\overline{\lambda\chi}, \overline{\lambda\chi'}}.$$

Since the action of λ is by simultaneously permuting rows and columns each $\lambda \cdot Y$, and hence Z itself, is positive semidefinite with ones on the diagonal. Since we have constructed Z by averaging over a group action, Z is fixed by the action and so satisfies $Z_{\overline{\lambda\chi}, \overline{\lambda\chi'}} = Z_{\chi, \chi'}$ for all $\chi, \chi' \in \widehat{G}$. Consequently there is some $(y_\chi)_{\chi \in \widehat{G}}$ such that $Z_{\chi, \chi'} = y_{\overline{\chi\chi'}}$ for all $\chi, \chi' \in \widehat{G}$. It remains to show that if $\overline{\chi\chi'} \in \mathcal{S}$ then $y_{\overline{\chi\chi'}} := Z_{\chi, \chi'} = Y_{\chi, \chi'} := \ell_{\overline{\chi\chi'}}$. This holds because if $\overline{\chi\chi'} \in \mathcal{S}$ then

$$Z_{\chi, \chi'} = \frac{1}{|\widehat{G}|} \sum_{\lambda \in \widehat{G}} Y_{\overline{\lambda\chi}, \overline{\lambda\chi'}} = \frac{1}{|\widehat{G}|} \sum_{\lambda \in \widehat{G}} \ell_{\overline{\lambda\chi\chi'}} = \ell_{\overline{\chi\chi'}}.$$

Hence y has all the desired properties, completing the proof. \square

Proof of Theorem 1D. First we show that $\mathcal{M}(G, \mathcal{S})$ is a subset of the right-hand-side of (22). To see this observe that the right-hand-side of (21) is certainly contained in the right-hand-side of (22).

We now establish the reverse inclusion. Suppose $(\ell_\chi)_{\chi \in \mathcal{S}}$ is such that there exists $(y_\chi)_{\chi \in \mathcal{T}^{-1}\mathcal{T}}$ with $y_\chi = \ell_\chi$ for all $\chi \in \mathcal{S}$, $y_{1_{\widehat{G}}} = 1$ and $M_{\mathcal{T}}(y) \succeq 0$. Let Γ be a chordal cover of $\text{Cay}(\widehat{G}, \mathcal{S})$ with Fourier support \mathcal{T} . Specifically Γ has the property that for every maximal clique \mathcal{C} of Γ there is $\chi_{\mathcal{C}} \in \widehat{G}$ such that $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$.

Define the Γ -partial matrix $Y_{\eta, \eta'} = y_{\overline{\eta\eta'}}$ whenever (η, η') form an edge of Γ and $Y_{\chi, \chi} = 1$ for all $\chi \in \widehat{G}$. This is well defined because if (η, η') is an edge of Γ , then $\overline{\eta\eta'} \in \mathcal{T}^{-1}\mathcal{T}$. To see this observe that any edge of Γ is contained in a maximal clique \mathcal{C} , and so there is some $\chi_{\mathcal{C}} \in \widehat{G}$ such that $\chi_{\mathcal{C}}\eta \in \mathcal{T}$ and $\chi_{\mathcal{C}}\eta' \in \mathcal{T}$. Consequently $\overline{\eta\eta'} = \overline{\chi_{\mathcal{C}}\eta\chi_{\mathcal{C}}\eta'} \in \mathcal{T}^{-1}\mathcal{T}$.

We show that $Y[\mathcal{C}, \mathcal{C}] \succeq 0$ for all maximal cliques \mathcal{C} of the chordal graph Γ . This holds because

$$Y[\mathcal{C}, \mathcal{C}] = [y_{\overline{\eta\eta'}}]_{\eta, \eta' \in \mathcal{C}} = [y_{\overline{\chi_{\mathcal{C}}\eta\chi_{\mathcal{C}}\eta'}}]_{\eta, \eta' \in \mathcal{C}} = [y_{\overline{\chi\chi'}}]_{\chi, \chi' \in \chi_{\mathcal{C}}\mathcal{C}}$$

which, since $\chi_{\mathcal{C}}\mathcal{C} \subseteq \mathcal{T}$, is a principal submatrix of the positive semidefinite (by assumption) matrix $M_{\mathcal{T}}(y) = [y_{\overline{\chi\chi'}}]_{\chi, \chi' \in \mathcal{T}}$. By the chordal completion theorem (Theorem 6) we can complete Y to a positive semidefinite matrix $Y \in \mathbf{S}_+^{\widehat{G}}$. The completed matrix has unit diagonal and, whenever $\overline{\chi\chi'} \in \mathcal{S}$,

$$Y_{\chi, \chi'} = y_{\overline{\chi\chi'}} = \ell_{\overline{\chi\chi'}}$$

where the first equality is because the edge set of Γ contains the edge set of $\text{Cay}(\widehat{G}, \mathcal{S})$ and the second is from the definition of y . Hence, by Proposition 5, $(\ell_\chi)_{\chi \in \mathcal{S}} \in \mathcal{M}(G, \mathcal{S})$, as we require. \square

A.2 Proof of Lemma 1

Proof of Lemma 1. First note that $L \in \mathbf{H}_+^d$ if and only if $\bar{L} \in \mathbf{H}_+^d$ which holds if and only if the block diagonal matrix $\begin{bmatrix} L & 0 \\ 0 & \bar{L} \end{bmatrix} \in \mathbf{H}_+^{2d}$. Conjugating by a unitary matrix we obtain

$$\begin{bmatrix} \frac{1}{\sqrt{2}}I & \frac{1}{\sqrt{2}}I \\ \frac{i}{\sqrt{2}}I & -\frac{i}{\sqrt{2}}I \end{bmatrix} \begin{bmatrix} L & 0 \\ 0 & \bar{L} \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}}I & \frac{1}{\sqrt{2}}I \\ \frac{i}{\sqrt{2}}I & -\frac{i}{\sqrt{2}}I \end{bmatrix}^* = \begin{bmatrix} \operatorname{Re}[L] & \operatorname{Im}[L] \\ -\operatorname{Im}[L] & \operatorname{Re}[L] \end{bmatrix}. \quad (31)$$

We have simply recovered the familiar realization of \mathbf{H}_+^d as a section of \mathbf{S}_+^{2d} , and have not yet used any special properties of \mathcal{L} . To complete the proof it remains to carefully choose a $2d \times 2d$ orthogonal matrix Q (depending on J) such that

$$Q \begin{bmatrix} \operatorname{Re}[L] & \operatorname{Im}[L] \\ -\operatorname{Im}[L] & \operatorname{Re}[L] \end{bmatrix} Q^T = \begin{bmatrix} \operatorname{Re}[L] - J \operatorname{Im}[L] & 0 \\ 0 & \operatorname{Re}[L] - J \operatorname{Im}[L] \end{bmatrix} \quad \text{for all } L \in \mathcal{L}.$$

Observe that $J^2 = I$ and $J^T J = I$ imply that $J = J^T$. Since $JLJ^T = \bar{L}$ we have that for all $L \in \mathcal{L}$,

$$\operatorname{Re}[L] = \frac{L + J L J}{2} \quad \text{and} \quad \operatorname{Im}[L] = \frac{L - J L J}{2i}. \quad (32)$$

It follows that for all $L \in \mathcal{L}$, $\operatorname{Re}[L]$ and $\operatorname{Im}[L]$ commute and anti-commute respectively with J , i.e.,

$$J \operatorname{Re}[L] = \operatorname{Re}[L] J \quad \text{and} \quad J \operatorname{Im}[L] = -\operatorname{Im}[L] J. \quad (33)$$

Choosing Q to be the orthogonal matrix $Q = \frac{1}{\sqrt{2}} \begin{bmatrix} I & J \\ -J & I \end{bmatrix}$ we obtain

$$\begin{bmatrix} \frac{1}{\sqrt{2}}I & \frac{1}{\sqrt{2}}J \\ -\frac{1}{\sqrt{2}}J & \frac{1}{\sqrt{2}}I \end{bmatrix} \begin{bmatrix} \operatorname{Re}[L] & \operatorname{Im}[L] \\ -\operatorname{Im}[L] & \operatorname{Re}[L] \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}}I & \frac{1}{\sqrt{2}}J \\ -\frac{1}{\sqrt{2}}J & \frac{1}{\sqrt{2}}I \end{bmatrix}^T = \begin{bmatrix} \operatorname{Re}[L] - J \operatorname{Im}[L] & 0 \\ 0 & \operatorname{Re}[L] - J \operatorname{Im}[L] \end{bmatrix}.$$

Clearly this last matrix is positive semidefinite if and only if the real symmetric matrix $\operatorname{Re}[L] - J \operatorname{Im}[L]$ is positive semidefinite, completing the proof. \square

A.3 Proof of Proposition 6

Proof of Proposition 6. The proof proceeds as follows. First we define a graph Γ and prove that it is a chordal cover of $\operatorname{Cay}(\widehat{G}, \mathcal{S})$. We then characterize the maximal cliques of Γ . Finally we show that for any maximal clique \mathcal{C} of Γ there is some $S \in 2^{[n]}$ such that $S \triangle \mathcal{C} \subseteq \mathcal{T}$, establishing the stated result. We consider the two cases $\lceil n/2 \rceil$ even and $\lceil n/2 \rceil$ odd separately. We describe the argument in detail in the case where $\lceil n/2 \rceil$ is even, and just sketch the required modifications in the case where $\lceil n/2 \rceil$ is odd.

Assume that $\lceil n/2 \rceil$ is even. Let Γ be the graph with vertex set $2^{[n]}$ such that two vertices S, T are adjacent in Γ if and only if either

- $|S|$ and $|T|$ are both even and $||S| - |T|| \leq 2$ or
- $|S|$ and $|T|$ are both odd and $||\phi(S)| - \phi(T)|| \leq 2$.

Note that just like $\operatorname{Cay}(\widehat{G}, \mathcal{S})$, the graph Γ also has two connected components with vertex sets $\mathcal{T}_{\text{even}}$ and \mathcal{T}_{odd} . Furthermore, ϕ (defined in Section 4.2) is also an automorphism of Γ that exchanges these two connected components. Observe that if $|S \triangle T| = 2$ (i.e. S and T are adjacent in $\operatorname{Cay}(\widehat{G}, \mathcal{S})$) then both $||S| - |T|| \leq 2$ and $||\phi(S)| - \phi(T)|| \leq 2$ hold. Hence if S and T are adjacent in $\operatorname{Cay}(\widehat{G}, \mathcal{S})$ they are also adjacent in Γ .

We now show that Γ is a chordal graph. Let the vertices $S_1, S_2, S_3, \dots, S_k$ form a k -cycle (with $k \geq 4$) in Γ such that each of the $S_i \in \mathcal{T}_{\text{even}}$. Without loss of generality assume that $|S_1| \leq |S_i|$ for $1 \leq i \leq k$. We show that the cycle $S_1, S_2, S_3, \dots, S_k$ has a chord. If $|S_2| = |S_1|$ then $||S_1| - |S_3|| = ||S_2| - |S_3|| \leq 2$ (since S_2 and S_3 are adjacent) and so there is a chord between S_1 and S_3 . Otherwise suppose $|S_2| = |S_1| + 2$. Because S_1 and S_k are adjacent we see that either $|S_k| = |S_1| = |S_2| - 2$ or $|S_k| = |S_1| + 2 = |S_2|$ and so there is a chord between S_2 and S_k . Now suppose $S_1, S_2, S_3, \dots, S_k$ form a k -cycle (with $k \geq 4$) in Γ such that each

of the $S_i \in \mathcal{T}_{\text{odd}}$. Then the image of the cycle under ϕ is a k -cycle in Γ with vertices in $\mathcal{T}_{\text{even}}$ and so it has a chord. Since ϕ is an automorphism of Γ it follows that $S_1, S_2, S_3, \dots, S_k$ also has a chord. So Γ is a chordal cover of $\text{Cay}(\widehat{G}, \mathcal{S})$.

The subgraphs of Γ induced by the vertex sets $\mathcal{C}_k := \mathcal{T}_k \cup \mathcal{T}_{k+2}$ (for $k = 0, 2, \dots, 2\lfloor n/2 \rfloor - 2$) and the vertex sets $\phi(\mathcal{C}_k)$ (for $k = 0, 2, \dots, 2\lfloor n/2 \rfloor - 2$) are cliques in Γ . In fact, these are maximal cliques in Γ . To show that each \mathcal{C}_k is a maximal clique, suppose S is a vertex that is not in \mathcal{C}_k . Then either $|S|$ is odd (in which case S is not adjacent to any element of \mathcal{C}_k) or $|S| \leq k - 2$ (in which case S is not adjacent to any $T \in \mathcal{T}_{k+2}$) or $|S| \geq k + 4$ (in which case S is not adjacent to any $T \in \mathcal{T}_k$). Hence there is no inclusion-wise larger clique of Γ containing \mathcal{C}_k . Since ϕ is an automorphism of Γ it follows that the $\phi(\mathcal{C}_k)$ are also maximal cliques of Γ . Finally, there are no other maximal cliques in Γ because every edge of Γ is contained either in \mathcal{C}_k or $\phi(\mathcal{C}_k)$ for some $k = 0, 2, \dots, 2\lfloor n/2 \rfloor - 2$.

It remains to show that for any maximal clique \mathcal{C}_k (for $k = 0, 2, \dots, 2\lfloor n/2 \rfloor - 2$) of Γ there is $S_k \in 2^{[n]}$ such that $S_k \triangle \mathcal{C}_k \subseteq \mathcal{T}$. This is sufficient to establish that $\text{Cay}(\widehat{G}, \mathcal{S})$ has a chordal cover with Fourier support \mathcal{T} because for the cliques $\phi(\mathcal{C}_k)$ we have that $\phi(S_k) \triangle \phi(\mathcal{C}_k) = S_k \triangle \mathcal{C}_k \subseteq \mathcal{T}$. The following gives valid choices of S_k (for $k = 0, 2, \dots, 2\lfloor n/2 \rfloor - 2$).

- If $k \leq \lfloor n/2 \rfloor - 2$ then $\mathcal{C}_k \subseteq \mathcal{T}$ so we can take $S_k = \emptyset$.
- If $k \geq \lfloor n/2 \rfloor$ and n is even then $n = 2\lfloor n/2 \rfloor$ and so $n - k - 2 \leq \lfloor n/2 \rfloor - 2$. Hence $[n] \triangle \mathcal{C}_k = \mathcal{C}_{n-k-2} \subseteq \mathcal{T}$ so we can take $S_k = [n]$.
- If $k \geq \lfloor n/2 \rfloor$ and n is odd then $n = 2\lfloor n/2 \rfloor + 1$ and so $n - k + 1 \leq \lfloor n/2 \rfloor$. Hence

$$\phi([n]) \triangle \mathcal{C}_k = [n] \triangle \phi(\mathcal{C}_k) \subseteq [n] \triangle (\mathcal{T}_{k-1} \cup \mathcal{T}_{k+1} \cup \mathcal{T}_{k+3}) \subseteq \mathcal{T}_{n-k-3} \cup \mathcal{T}_{n-k-1} \cup \mathcal{T}_{n-k+1} \subseteq \mathcal{T}$$

so we can take $S_k = \phi([n])$.

This completes the argument in the case where $\lfloor n/2 \rfloor$ is even.

In the case where $\lfloor n/2 \rfloor$ is odd we exchange the roles of the odd and even components in the definition of Γ and throughout the argument. More precisely, two vertices S, T are adjacent in Γ if and only if either

- $|S|$ and $|T|$ are both odd and $||S| - |T|| \leq 2$ or
- $|S|$ and $|T|$ are both even and $||\phi(S)| - |\phi(T)|| \leq 2$.

It is still the case that Γ is a chordal cover of $\text{Cay}(\widehat{G}, \mathcal{S})$. Its maximal cliques are now $\mathcal{C}_k := \mathcal{T}_k \cup \mathcal{T}_{k+2}$ for $k = 1, 3, \dots, 2\lfloor n/2 \rfloor - 3$ together with the $\phi(\mathcal{C}_k)$. Note that the cliques are now indexed by odd integers. As before, we can choose the S_k (for $k = 1, 3, \dots, 2\lfloor n/2 \rfloor - 3$) to be $S_k = \emptyset$ if $k \leq \lfloor n/2 \rfloor - 2$, $S_k = [n]$ if $k \geq \lfloor n/2 \rfloor$ and n is even, and $S_k = \phi([n])$ if $k \geq \lfloor n/2 \rfloor$ and n is odd.

This completes the argument in the case where $\lfloor n/2 \rfloor$ is odd. \square

A.4 Proof of Theorem 9: triangulation of the cycle graph

In this appendix we prove Theorem 9 concerning the triangulation of the cycle graph C_N . Theorem 10 below shows how to construct a triangulation of the cycle graph C_{N+1} on $N + 1$ nodes, by induction. The triangulation of C_N used to obtain Theorem 9 will then be obtained simply by contracting a certain edge of the triangulation of C_{N+1} (more details below). We thus start by describing a triangulation of the $N + 1$ -cycle.

Theorem 10 (Triangulation of the cycle graph on $N + 1$ vertices). *Let N be an integer greater than or equal to 2. Let $k_1 < \dots < k_l$ be the position of the nonzero digits in the binary expansion of N , i.e., $N = \sum_{i=1}^l 2^{k_i}$. Let k be the largest integer such that $2^k < N$ (i.e., $k = k_1 - 1$ if N is a power of two and $k = k_l$ otherwise). Then there exists a triangulation of the cycle graph on $N + 1$ nodes C_{N+1} with frequencies:*

$$\mathcal{T} = \{0\} \cup \{\pm 2^i, i = 0, \dots, k\} \cup \left\{ \sum_{j=1}^i 2^{k_j}, i = 1, \dots, l - 1 \right\}. \quad (34)$$

Proof. The proof of the theorem is by induction on N . Consider the cycle graph on $N + 1$ nodes where nodes are labeled $0, 1, \dots, N$. To triangulate the graph, we first put an edge between nodes 0 and 2^k and another edge between nodes 2^k and N , where 2^k is the largest power of two that is strictly smaller than N . This is depicted in Figure 7.




Figure 7. Recursive triangulation of the cycle $0 \dots N$ on $N + 1$ vertices

Note that the frequencies used by the triangle $\{0, 2^k, N\}$ are equivalent, by translation, to $\{-2^k, 0, N - 2^k\}$. We now use induction to triangulate the two remaining parts of the cycle (denoted (a) and (b) in Figure 7):

- For part (a), which is a cycle graph labeled $0 \dots N'$ with $N' = 2^k$, the induction hypothesis gives us a triangulation with frequencies

$$\mathcal{T}_a = \{0\} \cup \{\pm 2^i, i = 0, \dots, k - 1\}. \quad (35)$$

- For part (b) of the graph, we use induction on the cycle $2^k \dots N$ which is, by translation, equivalent to the cycle with labels $0 \dots N''$ where $N'' = N - 2^k$. We distinguish two cases:

- If $N = 2^{k+1}$, then we have $N'' = 2^k$ and induction gives a triangulation of (b) with the same frequencies as for part (a). Thus in this case we get a triangulation of the full $(N + 1)$ -cycle with frequencies:

$$\mathcal{T}_a \cup \{-2^k, 0, 2^k\} = \{0\} \cup \{\pm 2^i, i = 0, \dots, k\}$$

which is what we want.

- Now assume that $N < 2^{k+1}$, which means that the most significant bit of N is at position $k = k_l$. Thus the binary expansion of $N'' = N - 2^k$ is the same as that of N except that the bit at position $k = k_l$ is replaced with a 0. Let k'' be the largest integer such that $2^{k''} < N''$. Using induction we get a triangulation of the cycle $0 \dots N''$ using frequencies where

$$\mathcal{T}_b = \{0\} \cup \{\pm 2^i, i = 0, \dots, k''\} \cup \left\{ \sum_{j=1}^i 2^{k_j}, j = 1, \dots, l - 2 \right\}. \quad (36)$$

Combining the triangulation of parts (a) and part (b) we get a triangulation of the $(N + 1)$ -cycle with frequencies

$$\underbrace{\{-2^k, 0, N - 2^k\}}_{\text{triangle } \{0, 2^k, N\}} \cup \mathcal{T}_a \cup \mathcal{T}_b.$$

Given the expressions (35) and (36) for \mathcal{K}_a and \mathcal{K}_b , and noting that $k'' \leq k - 1$ and that $N - 2^k = \sum_{j=1}^{l-1} 2^{k_j}$, one can check that the triangulation has frequencies in

$$\mathcal{T} = \{0\} \cup \{\pm 2^i, i = 0, \dots, k\} \cup \left\{ \sum_{j=1}^i 2^{k_j}, i = 1, \dots, l - 1 \right\}.$$

which is exactly what we want.

To complete the proof, it remains to show the base case of the induction. We will show the base cases $N = 2$ and $N = 3$. For $N = 2$, note that the $(N + 1)$ -cycle is simply a triangle which is already triangulated

and the frequencies are simply $\{-1, 0, 1\}$. If we evaluate expression (34) for $N = 2$ (note that here $k = 0$) we get $\mathcal{T} = \{-1, 0, 1\}$, as needed.

For $N = 3$ (the 4-cycle), we have $k = 1$ and $l = 2$ with $k_1 = 0$ and $k_2 = 1$. Thus expression (34) evaluates to $\mathcal{T} = \{0\} \cup \{\pm 1, \pm 2\} \cup \{1\} = \{-2, -1, 0, 1, 2\}$. It is easy to construct a triangulation of the 4-cycle with such frequencies (one can even construct one where $\mathcal{T} = \mathcal{K} \cup (-\mathcal{K}) = \{-1, 0, 1\}$). \square

Example 8. Figure 8 shows the recursive construction for the case $N = 8$. We have indicated in each triangle (3-clique) the associated set of frequencies. \diamond




Figure 8. Illustration of the recursive triangulation of the $(N + 1)$ -cycle for $N = 8$.

Proof of Theorem 9. To prove Theorem 9 for the N -cycle, we use the triangulation of the $(N + 1)$ -cycle of Theorem 10 except that we regard nodes 0 and N as the same nodes (they collapse into a single one). Thus this means that the triangle in Figure 7 with frequencies $\{-2^k, 0, N - 2^k\}$ also collapses and we only have to look at the frequencies for parts (a) and (b). It is not hard to show that the frequencies we get are the same as those given in Equation (34) except that in the middle term the iterate i goes from 0 to $k - 1$ (instead of from 0 to k), and in the last term the iterate i goes from 1 to $l - 2$ (instead of from 1 to $l - 1$) which gives exactly the set of frequencies of Equation (26). \square

Note that there are actually many different ways of constructing triangulations for the cycle graph, and different constructions will lead to a different set of “frequencies”. We can mention that for the cycle graph C_N one can actually construct a triangulation where the number of frequencies is related to the logarithm of N base 3. When N is a power of three the frequencies are precisely the powers of 3 that are smaller than N . We omit the precise description of this construction, but Figure 9 shows the triangulation for the 9-cycle and 27-cycle.




Figure 9. Triangulation of the 9-cycle with frequencies $\mathcal{T} = \{0, \pm 1, \pm 3\}$ and of the 27-cycle with frequencies $\mathcal{T} = \{0, \pm 1, \pm 3, \pm 9\}$.

References

- [AHMR88] Jim Agler, William Helton, Scott McCullough, and Leiba Rodman. Positive semidefinite matrices with a given sparsity pattern. *Linear algebra and its applications*, 107:101–149, 1988. [8](#)
- [Bar02] Alexander Barvinok. *A course in convexity*, volume 54. American Mathematical Society, 2002. [4](#)
- [BFMP14] Yuri Bogomolov, Samuel Fiorini, Aleksandr Maksimenko, and Kanstantsin Pashkovich. Small extended formulations for cyclic polytopes. *arXiv preprint arXiv:1401.8138*, 2014. [5](#), [25](#)
- [BGP14] Grigoriy Blekherman, João Gouveia, and James Pfeiffer. Sums of squares on the hypercube. *arXiv preprint arXiv:1402.4199*, 2014. [3](#)
- [BPT13] Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas. *Semidefinite optimization and convex algebraic geometry*. SIAM, 2013. [2](#)
- [Dum07] Bogdan Dumitrescu. *Positive trigonometric polynomials and signal processing applications*. Springer, 2007. [4](#)
- [FGP⁺14] Hamza Fawzi, João Gouveia, Pablo A. Parrilo, Richard Z. Robinson, and Rekha R. Thomas. Positive semidefinite rank. *arXiv preprint arXiv:1407.4095*, 2014. [26](#)
- [FKPT13] Samuel Fiorini, Volker Kaibel, Kanstantsin Pashkovich, and Dirk Oliver Theis. Combinatorial bounds on nonnegative rank and extended formulations. *Discrete Mathematics*, 313(1):67 – 83, 2013. [5](#), [25](#)
- [FSP13] Hamza Fawzi, James Saunderson, and Pablo A. Parrilo. Equivariant semidefinite lifts and sum-of-squares hierarchies. *arXiv preprint arXiv:1312.6662*, 2013. [26](#)
- [FSP14] Hamza Fawzi, James Saunderson, and Pablo A. Parrilo. Equivariant semidefinite lifts of regular polygons. *arXiv preprint arXiv:1409.4379*, 2014. [5](#), [21](#), [26](#)
- [Gal63] David Gale. Neighborly and cyclic polytopes. In *Proceedings of the Seventh Symposium in Pure Mathematics of the American Mathematical Society*, volume 7, pages 225–232, 1963. [5](#), [25](#)
- [GJSW84] Robert Grone, Charles R. Johnson, Eduardo M. Sá, and Henry Wolkowicz. Positive definite completions of partial hermitian matrices. *Linear algebra and its applications*, 58:109–124, 1984. [5](#), [6](#), [8](#), [9](#)
- [GPT13] João Gouveia, Pablo A. Parrilo, and Rekha R. Thomas. Lifts of convex sets and cone factorizations. *Mathematics of Operations Research*, 38(2):248–264, 2013. [9](#)
- [GT84] Andreas Griewank and Philippe L. Toint. On the existence of convex decompositions of partially separable functions. *Mathematical Programming*, 28(1):25–49, 1984. [5](#), [6](#), [8](#)
- [KP11] Volker Kaibel and Kanstantsin Pashkovich. Constructing extended formulations from reflection relations. In *Integer Programming and Combinatorial Optimization*, pages 287–300. Springer, 2011. [25](#)
- [Las01] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. [10](#)
- [Lau03] Monique Laurent. Lower bound for the number of iterations in semidefinite hierarchies for the cut polytope. *Mathematics of Operations Research*, 28(4):871–883, 2003. [3](#), [4](#), [5](#), [18](#), [20](#), [26](#)
- [Nes00] Yurii Nesterov. Squared functional systems and optimization problems. In *High performance optimization*, pages 405–440. Springer, 2000. [10](#)
- [Par00] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, 2000. [10](#)

- [Rud90] Walter Rudin. *Fourier Analysis on Groups*. John Wiley & Sons, Inc., 1990. [6](#), [14](#)
- [Ter99] Audrey Terras. *Fourier analysis on finite groups and applications*. Cambridge University Press, 1999. [6](#)
- [Yan91] Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991. [9](#)
- [Zie95] Günter M. Ziegler. *Lectures on polytopes*, volume 152. Springer, 1995. [4](#), [25](#)